Finding the Maximum Subset with Bounded Convex Curvature*

Mikkel Abrahamsen[†1] **and Mikkel Thorup**[‡2]

1,2 Department of Computer Science, University of Copenhagen
Universitetsparken 5
DK-2100 Copenhagen Ø
Denmark
miab@di.ku.dk, mikkel2thorup@gmail.com

arXiv:1603.02080v1 [cs.CG] 7 Mar 2016

Abstract

We describe an algorithm for solving an important geometric problem arising in computer-aided manufacturing. When machining a pocket in a solid piece of material such as steel using a rough tool in a milling machine, sharp convex corners of the pocket cannot be done properly, but have to be left for finer tools that are more expensive to use. We want to determine a tool path that maximizes the use of the rough tool. Mathematically, this boils down to the following problem. Given a simply-connected set of points P in the plane such that the boundary ∂P is a curvilinear polygon consisting of n line segments and circular arcs of arbitrary radii, compute the maximum subset $Q \subseteq P$ consisting of simply-connected sets where the boundary of each set is a curve with bounded convex curvature. A closed curve has bounded convex curvature if, when traversed in counterclockwise direction, it turns to the left with curvature at most 1. There is no bound on the curvature where it turns to the right. The difference in the requirement to left- and right-curvature is a natural consequence of different conditions when machining convex and concave areas of the pocket. We devise an algorithm to compute the unique maximum such set Q. The algorithm runs in $O(n \log n)$ time and uses $O(n)$ space.

For the correctness of our algorithm, we prove a new generalization of the Pestov-Ionin Theorem. This is needed to show that the output Q of our algorithm is indeed maximum in the sense that if Q' is any subset of P with a boundary of bounded convex curvature, then $Q' \subseteq Q$.

1 Introduction

The motivation for our work comes from the generation of toolpaths for pocket machining. Pocket machining is the process of cutting some specified pocket in a piece of material – in our case most likely a piece of metal – using a milling machine. We first describe the clean mathematical problem that we solve and afterwards explain how it relates to pocket machining.

Consider a simply-connected closed subset of the plane and the weakly simple, closed curve around its boundary which we traverse in counter-clockwise direction. We say that the curve is convex where it turns left and concave where it turns right. We say it has *bounded convex curvature* if it turns left with curvature at most 1. There is no bound on the right-curvature, and we may even have sharp concave corners. We say that a set in the plane has bounded convex curvature if all the connected components are simply-connected

* A short version of this paper will appear at SoCG 2016.

† Research partly supported by Mikkel Thorup's Advanced Grant from the Danish Council for Independent Research under the Sapere Aude research career programme.

‡ Research partly supported by an Advanced Grant from the Danish Council for Independent Research under the Sapere Aude research career programme.

and the weakly simple, closed curve around the boundary of each connected component has bounded convex curvature. Similarly, we say that a closed curve has *bounded curvature in general* if it turns to the left and to the right with curvature at most 1.

A nice composition property of sets of bounded convex curvature is that if we take two such sets Q_1 and Q_2, unite them, and fill out any holes, then the resulting set, denoted $Q_1 \uplus Q_2$, has bounded convex curvature. See Figure 1, where the boundary of $Q_1 \uplus Q_2$ is the thick curve. This composition property does not hold if we demand that the curvature is bounded in general.

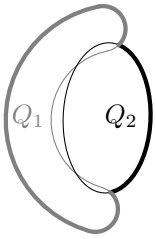

Figure 1

The input to our problem is a simply-connected closed set of points P in the plane such that the boundary ∂P is a curvilinear polygon consisting of n line segments and circular arcs of arbitrary radii. We present an algorithm that in $O(n \log n)$ time finds the unique maximum subset $Q \subseteq P$ of bounded convex curvature, that is, Q contains any other $Q' \subseteq P$ of bounded convex curvature. See Figure 2 for an example.

We note that the uniqueness of a maximal subset Q of bounded convex curvature follows from the composition property; for if there was another $Q' \subseteq P$ of bounded convex curvature that was not contained in Q, then $Q \uplus Q'$ would also be contained in P and have bounded convex curvature, contradicting the maximality of Q.

The boundary of Q will be a curvilinear polygon consisting of $O(n)$ line segments and circular arcs. A very useful property of the boundary of Q is that all concave arcs and vertices are also on the boundary of P. Indeed, it is easy to verify that if there is a concave arc or vertex on the boundary of Q which is not on the boundary of P, then Q is not maximal. A similar reasoning implies that if the boundary of P is a simple curve, then so is the boundary of Q.

We now describe different contexts in which this problem appears naturally. The general problem is that we are given an area S of the plane whose boundary is represented as a curvilinear polygon. There is a thin layer of material in S close to the boundary ∂S of S. The goal is to remove that layer without removing anything from outside S. We are given a rough tool, and we want to remove as much as possible of the thin layer, leaving as little as possible for finer tools that are more expensive to use. The output is a toolpath for the rough tool consisting of one or more curvilinear polygons. The tool is a disk $\mathscr{D}$ of radius r, where r is bigger than the width of the layer we wish to remove. The toolpath is the path which the center of $\mathscr{D}$ is following and the material cut away is the area swept by $\mathscr{D}$ as its center follows the toolpath. The reason that we only have to handle a thin layer close to the boundary of S is because the area farther from the boundary is removed beforehand by tools that are less precise since they do not get close to the boundary. Thus we may assume that the material at all points at some distance $\delta \leq r$ from the boundary have been removed. Some of the points closer to the boundary may also have been removed, but this only makes it easier for our tool to move. With this in mind, when the tool follows a weakly simple closed curve, we think of it as removing the interior of the curve plus every point at distance at most r to the curve.

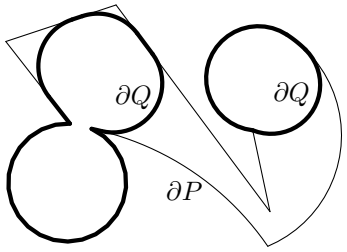

Figure 2

Let P be the inwards offset of S by r, that is, P is the subset of S of points with distance at least r to the pocket boundary ∂S. See Figure 3. P is the set of all allowed positions of the center of $\mathscr{D}$. If we had complete control over the tool, then we would be able to remove the material in all of P and at all points with distance at most r from P by letting the tool

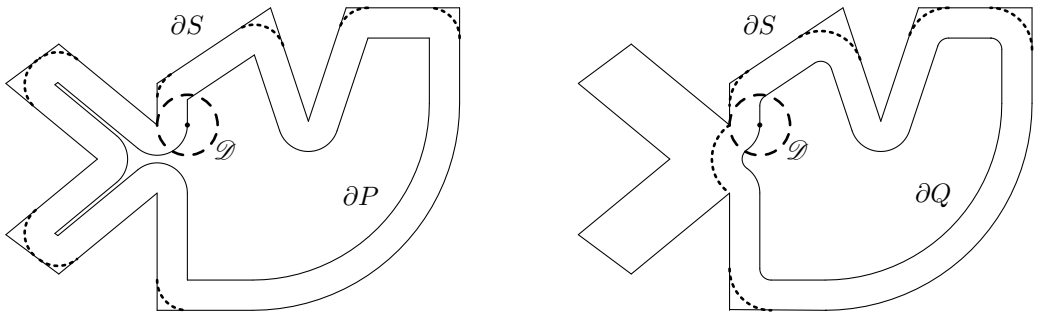

Figure 3 A pocket bounded by ∂S. To the left is shown the boundary ∂P of the inwards offset of S by r. To the right is shown the boundary ∂Q of the maximum subset with bounded convex curvature of P. The dotted arcs in the corners show the boundary of the material in S that cannot be removed by $\mathscr{D}$ using the two toolpaths.

center traverse the boundary ∂P. However, there are restrictions on what tool paths we can follow, e.g., we cannot count on following a toolpath with sharp corners in a precise way.

We are now ready to describe the first application where we want to compute the maximum subset of bounded convex curvature. Assume that the tool can only follow a path which has bounded curvature in general. Assume furthermore that the tool can turn at least as sharply as its own boundary, that is, $r \geq 1$.

Using our algorithm for bounded convex curvature, we are able to identify the maximum area that can be removed by the tool such that the toolpath has bounded curvature in general. First we compute the above set P which is the inwards offset of S by r. The boundary ∂P can be computed from the Voronoi diagram of ∂S [8]. Clearly the toolpath has to stay inside P. We now note that every concave part of ∂P has curvature at most $1/r \leq 1$. Next we use our algorithm to find the maximum subset $Q \subseteq P$ of bounded convex curvature, see Figure 3. The area cut away as $\mathscr{D}$ follows ∂Q is the unique maximum subset that can be cut out of S using a tool with radius $r \geq 1$ and such that the toolpath has bounded convex curvature. However, all concave arcs and vertices on ∂Q stem from P which has concave curvature at most $1/r \leq 1$. It follows that using the toolpath ∂Q, we cut out the maximum subset of S under the condition that the toolpath has bounded curvature in general. This also implies that we remove the maximum subset of the thin layer of material close to the pocket boundary ∂S.

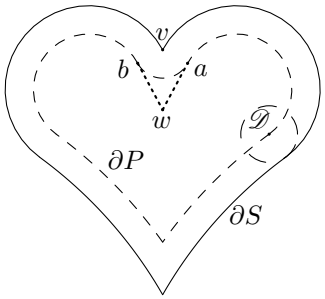

Figure 4

In the above example, the set P had bounded concave curvature. In particular, all concave arcs on ∂P have radius at least r, and for each concave arc A on ∂P of radius r and center v, there is an associated concave vertex v of ∂S. Let a and b be the first and last point on A. When the tool follows A, the corner v will be on the tool boundary $\partial\mathscr{D}$, and the slightest imprecision will blunt the corner v. A recommended alternative [12] is that we substitute A with two line segments aw and wb tangential to A at a and b, respectively, thus creating a sharp concave corner w on the toolpath, see Figure 4. Using this technique, the corner v will be cut much sharper and more precise. One can think of various variations of this technique, since we can "casually" stop and turn the tool at any point on its way to w because the remaining toolpath already ensures that all material will be cut away. This

shows that we cannot in general assume that there is any bound on the concave curvature of the input toolpath. We also note the asymmetry with convex corners and arcs, where overshooting a convex corner implies an illegal cut through the boundary of S.

We shall now provide a completely different explanation for the need for bounded convex curvature. The point is that it is often preferable for the surface quality of the resulting part that the tool moves with a constant speed. Recall that the tool is only removing a thin layer close to the pocket boundary. The width of this layer is typically a deliberately chosen small fraction of the tool radius r. When moving at constant speed, a convex turn implies a higher engagement of the tool in the sense of the amount of material removed per time unit. In concave turns the engagement is only decreased. A too high engagement could break the tool, and therefore we must bound the convex curvature of the tool path.

These and other issues related to the machining of corners have been extensively studied in the technical literature on pocket machining. See for instance the papers [6, 7, 10, 13, 17]. There are several previous papers suggesting methods to get bounded convex curvature, but none of them guarantees an optimal solution like ours. One idea for how to handle convex corners is to replace each of them by a convex circular arc as deep in the corner as possible. This is suggested and studied in the papers [6, 10, 13]. However, in all the papers it is assumed that every corner is formed by two line segments which are sufficiently long (relative to the angle between them) that a tangential corner-rounding arc of sufficient size can be placed inside the wedge they form. As can be seen in Figure 3, this is not always the case, and rounding a toolpath can require more complicated modifications.

One heuristic used to obtain a non-trivial subset of bounded convex curvature is the *double offset method*, where we offset P inwards by 1 and then offset the result outwards by 1 and use that as Q. This can be computed in $O(n \log n)$ time using Voronoi diagrams [16]. However, the method does not in general result in the maximum subset of bounded convex curvature. See for instance Figure 5, where P has bounded convex curvature and all material in S will be machined using ∂P as the toolpath. ∂S is the thick solid curve and ∂P is the thin solid curve. The innermost dotted curves are the boundary of the inwards offset of P by 1. The area of S between the thick dashed arcs will not be machined if the double offset method is used.

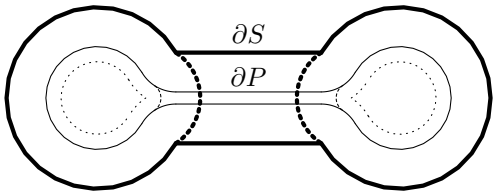

Figure 5

We establish the mathematical foundation of our work in Section 2 by introducing a new generalization of the theorem of Pestov and Ionin. In Section 3 we describe an algorithm that computes the maximum subset of bounded convex curvature of P. In Section 4 we describe how to implement the algorithm so that it uses $O(n)$ space and runs in $O(n \log n)$ time.

1.1 General notation and conventions

If M is a set of points in the plane, ∂M denotes the boundary of M and $\overline{M}$ denotes the closure of M. Let γ be a simple curve and x and y two points on γ. Then $\gamma[x, y]$ is the portion of γ between x and y including x and y. Similarly, $\gamma(x, y)$, $\gamma[x, y)$, and $\gamma(x, y]$ are used in the obvious way when none or one of x, y is included. If γ is a closed curve, $\gamma[x, y]$ denotes the portion of γ from x to y in the counterclockwise direction. As before, round parenthesis can be used to exclude one or both endpoints. If γ is not closed, the order of x and y does not matter. We say that a point $z \in \gamma$ is an *inner point* of γ if z is not an endpoint of γ.

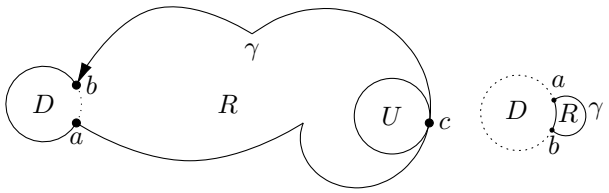

Figure 6 Cases where Theorem 1 does and does not apply.

2 Mathematical foundation: The theorem of Pestov and Ionin

An oriented, simple curve γ is in the class $\mathcal{L}^+$ if γ turns to the left with curvature at most 1, but γ is allowed to turn to the right arbitrarily sharply. There can be corners on γ where γ is not differentiable, but in such corners γ has to turn to the right. These notions are defined rigorously below. Note that a portion of a closed curve with bounded convex curvature traversed in counterclockwise direction is a curve in $\mathcal{L}^+$.

The correctness of the algorithm presented in this paper depends on the following generalization of the Pestov-Ionin Theorem, see Figure 6.

▶ **Theorem 1.** *Consider an open disk D of arbitrary radius and a curve γ in $\mathcal{L}^+$ from a to b such that $\gamma \cap \overline{D} = \{a, b\}$. Let R be the region bounded by $\partial D[b, a]$ and γ. If R contains D, then R contains an open unit disk U such that there exists a point $c \in \partial U \cap \gamma(a, b)$.*

From the theorem, the following corollary easily follows.

▶ **Corollary 2.** *Every simple, closed curve γ with bounded convex curvature contains an open unit-disk in its interior.*

The original theorem by Pestov and Ionin [14] was similar to Corollary 2, but γ was assumed to have bounded curvature in general. Howard and Treibergs [9] proved Corollary 2 for closed curves γ with bounded convex curvature for a more restricted class of curves. Ahn et al. [1] proved a theorem similar to Theorem 1, but assuming that γ has bounded curvature in general. Hence, the version of the theorem presented here generalizes all previous versions known to the authors of the present paper. The reader who is not interested in the formal definitions and the proof of the theorem can jump directly to Section 3.

2.1 Definition of the class $\mathcal{L}^+$ and proof of Theorem 1

Let $\gamma : [0, L] \to \mathbb{R}^2$ be a simple, continuously differentiable curve parameterized by arc length. We sometimes use γ as a short-hand notation for the set of points $\gamma([0, L])$ on γ. We define the coordinate fuctions γ_x and γ_y of γ such that $\gamma(s) = (\gamma_x(s), \gamma_y(s))$. Let $\mathbf{n}(s) = \widehat{\gamma'}(s) = (-\gamma'_y(s), \gamma'_x(s))$ be the unit normal of γ in s.

We define $\mathcal{K}^+$ to be the simple, continuously differentiable curves $\gamma : [0, L] \to \mathbb{R}^2$ parameterized by arc length such that for every $s \in [0, L)$ there exists a $H > 0$ such that for all $h \in [0, H]$ we have $\langle \gamma'(s + h) - \gamma'(s), \mathbf{n}(s) \rangle \leq h$, where $\langle \cdot, \cdot \rangle$ is the dot-product. We denote the class of all such curves as $\mathcal{K}^+$. Howard and Treibergs [9] use a similar definition. We first prove a theorem similar to Theorem 1, but for curves in $\mathcal{K}^+$ which is a subset of $\mathcal{L}^+$. In Section 2.1.1, we prove the general result by showing how to handle curves with sharp corners.

▶ **Example 3.** *Consider a counterclockwise arc of radius $r > 0$, $\gamma(s) = (r\cos\frac{s}{r}, r\sin\frac{s}{r})$. We only need to consider $s = 0$ to check if γ is in $\mathcal{K}^+$. We get $\langle\gamma'(h) - \gamma'(0), \mathbf{n}(0)\rangle = \sin\frac{h}{r}$. Since $\lim_{h\to 0}\frac{\sin\frac{h}{r}}{h/r} = 1$, γ is in $\mathcal{K}^+$ if and only if $r \geq 1$.*

On the other hand, if $\gamma(s) = (r\cos\frac{s}{r}, -r\sin\frac{s}{r})$ is a clockwise arc, we get $\langle\gamma'(h) - \gamma'(0), \mathbf{n}(0)\rangle = -\sin\frac{h}{r} < h$, so γ is in $\mathcal{K}^+$ for any r.

Whenever γ is continuously differentiable and parameterized by arc length, there exists a continuous function $\vartheta : [0, L] \to \mathbb{R}$ such that

$$\gamma'(s) = (\cos\vartheta(s), \sin\vartheta(s))$$

and

$$\gamma(s) = \gamma(0) + \left(\int_0^s \cos\vartheta(r)\,dr, \int_0^s \sin\vartheta(r)\,dr\right).$$

We denote ϑ as the *tangent indicatrix* associated to γ. When more curves are in play, we denote the tangent indicatrix of γ as ϑ_γ.

If p is a point and S is a set of points, $\mathrm{dist}(p, S) = \inf_{s\in S}\|p - s\|$. We usually define distances to compact sets S so that there is a point $s \in S$ for which $\mathrm{dist}(p, S) = \|p - s\|$.

▶ **Lemma 4.** *Let $s > 0$ and $f : [0, s] \to \mathbb{R}$ be a continuous function such that $f(0) = 0$ and $\vartheta(s) = (1+\varepsilon)\cdot s$ for $\varepsilon > 0$. Then there exists an $r \in [0, s)$ such that for all $x \in [r, s]$ we have $f(x) \geq f(r) + (1+\varepsilon)(x - r)$.*

Proof. Consider the function $g(x) = f(x) - (1+\varepsilon)\cdot x$. Choose a global minimum r of g. If g is a non-negative function, we can choose $r = 0$. Therefore, $r < s$ in any case. For $x \in [r, s]$, we now have $f(x) - (1+\varepsilon)\cdot x \geq f(r) - (1+\varepsilon)\cdot r$ and hence $f(x) \geq f(r) + (1+\varepsilon)(x - r)$. ◀

▶ **Lemma 5.** *Let $\gamma : [0, L] \to \mathbb{R}^2$ be a continuously differentiable curve parameterized by arc length and let $\vartheta : [0, L] \to \mathbb{R}$ be the associated tangent indicatrix. Then the following two conditions are equivalent:*

1. *γ is in $\mathcal{K}^+$.*
2. *For all $p \in [0, L)$ and for all $q \in [p, L]$ we have $\vartheta(q) - \vartheta(p) \leq q - p$.*

Proof. Assume first that condition 2 holds. Without loss of generality, we may assume that $\vartheta(0) = 0$ and prove that γ is in $\mathcal{K}^+$ by verifying the curvature condition for $s = 0$. Choose $H \in (0, \pi/2]$ such that $\vartheta([0, H]) \subseteq [-\pi/2, \pi/2]$. For every $h \in [0, H]$ we now have $\langle\gamma'(h) - \gamma'(0), \mathbf{n}(0)\rangle = \sin\vartheta(h) \leq h$, since $\sin x$ is increasing for $x \in [-\pi/2, \pi/2]$ and $\vartheta(h) \leq h$ for all h. Therefore, the curvature condition is satisfied.

Assume now that condition 2 does not hold, i.e., that $\vartheta(q) - \vartheta(p) > q - p$ for some q and p, $p \leq q$. Lemma 4 gives that there exists an $r \in [p, q)$ such that for all $x \in [r, q]$ we have $\vartheta(x) \geq \vartheta(r) + (1+\varepsilon)(x - r)$. Assume without loss of generality that $r = 0$ and $\vartheta(0) = 0$ so that $\vartheta(x) \geq (1+\varepsilon)\cdot x$ for all $x \in [0, q]$. Then $\langle\gamma'(x) - \gamma'(0), \mathbf{n}(0)\rangle = \sin\vartheta(x) > x$ when x is small enough and the curvature condition is not satisfied. ◀

A *wedge* is the intersection between two closed half-planes defined by two non-parallel lines. The *tip* of the wedge is the intersection point of the two lines defining the half-planes.

▶ **Lemma 6.** *Let $\gamma : [0, L] \to \mathbb{R}^2$ be a curve in $\mathcal{K}^+$ and let D be an open unit disk with center $\gamma(s) + \mathbf{n}(s)$ for some fixed $s \in [0, L]$. Then γ is disjoint from D in a neighborhood around $\gamma(s)$, i.e., there exists a $h > 0$ such that $\gamma(I) \cap D = \emptyset$ for $I = [s - h, s + h] \cap [0, L]$.*

Proof. We assume that $s < L$ and only show the existance of $h > 0$ such that $\gamma([s, s+h]) \cap D = \emptyset$. The other case follows from symmetry. Without loss of generality, assume that $\gamma(0) = (0, -1)$, $\gamma'(0) = (1, 0)$, $\vartheta(0) = 0$, and $s = 0$, so that D is the unit disk with center $(0, 0)$. Choose $h \in (0, \pi/4]$ such that $\vartheta([0, h]) \subseteq [-\pi/4, \pi/4]$. Lemma 5 gives that $\vartheta(x) \leq x$, and since $\cos x - \sin x$ is decreasing for $x \in [-\pi/4, \pi/4]$, we get for $x \in [0, h]$ that

$$\cos x - \sin x \leq \cos \vartheta(x) - \sin \vartheta(x).$$

By integrating over $x \in [0, s]$ for $s \in [0, h]$, we get

$$\cos s + \sin s \leq \gamma_x(s) - \gamma_y(s).$$

Since $\sin x$ is increasing for $x \in [-\pi/4, \pi/4]$, we also have

$$\gamma_y(s) = \gamma_y(0) + \int_0^s \sin \vartheta(x)\, dx \leq -1 + \int_0^s \sin x\, dx = -\cos s.$$

Therefore, for any $s \in [0, h]$, $\gamma(s)$ is in the area

$$A_s = \{(x, y) \mid y \leq x - \cos s - \sin s \text{ and } y \leq -\cos s\}.$$

A_s is the wedge of points between the two half-lines starting at $(\sin s, -\cos s)$ with directions $(-1, -1)$ and $(1, 0)$, respectively. Therefore, A_s and D are clearly disjoint when $0 \leq s \leq \pi/4$, and hence $\gamma([0, h])$ is disjoint from D. ◀

In the remaining part of this section, we consider an open disk D of arbitrary radius δ and a curve $\gamma : [0, L] \to \mathbb{R}^2$ in class $\mathcal{K}^+$ such that $\gamma \cap \overline{D} = \{\gamma(0), \gamma(L)\}$. The endpoints of γ divide ∂D into two arcs $\partial D[\gamma(0), \gamma(L)]$ and $\partial D[\gamma(L), \gamma(0)]$. One of them is interior and one is exterior in the sense that the interior is contained in the region bounded by the exterior and γ. Let R be the open region bounded by $\partial D[\gamma(L), \gamma(0)]$ and γ. Assume that $\partial D[\gamma(0), \gamma(L)] \subset \overline{R}$ is interior. Observe that R is to the left of γ and that $D \subset R$.

For the following lemma, it is actually not necessary that $\gamma \in \mathcal{K}^+$.

▶ **Lemma 7.** *Let g be an inner point on γ and let U be the largest open disk contained in R such that ∂U contains g. Then $\partial R \cap \partial U \subset \gamma$.*

Proof. Assume for contradiction that ∂U contains a point $d \in \partial D(\gamma(L), \gamma(0))$. This means that the radius of U is at most δ and that ∂U is contained in $\overline{D}$. Therefore, ∂U cannot pass through the point g for $g \notin \{\gamma(0), \gamma(L)\}$. ◀

Let $\overline{\gamma}$ be the closed curve obtained by adding the line segment from $\gamma(L)$ to $\gamma(0)$ to γ. $\overline{\gamma}$ is simple since the added line segment is in $\mathcal{D}$. Define the *tangent* of γ with angle φ as the line T_φ with normal vector $\mathbf{m}(\varphi) = (\cos\varphi, \sin\varphi)$ such that T_φ contains $\gamma(s')$ for $s' = \operatorname{argmax}_{s \in [0,L]} \langle \gamma(s), \mathbf{m}(\varphi) \rangle$. T_φ is a tangent of the convex hull of $\overline{\gamma}$ which contains the extreme points in direction $\mathbf{m}(\varphi)$.

Let $\mathcal{H}$ be the boundary of the convex hull of $\overline{\gamma}$. Since $\overline{\gamma}$ is a simple, closed curve, $\mathcal{H}$ is as well. The following Lemma is well-known.

▶ **Lemma 8.** *The cyclic order of the points on $\mathcal{H}$ is the same as on $\overline{\gamma}$, that is, let $a, b, c \in \mathcal{H} \cap \overline{\gamma}$ and assume that $b \in \mathcal{H}(a, c)$. Then $b \in \overline{\gamma}(a, c)$.*

▶ **Lemma 9.** *If $L < \pi$, there is no angle φ and $a, b \in [0, L]$ such that $\gamma(a) \in T_\varphi$, $\gamma(b) \in T_{\varphi+\pi}$, and $\gamma'(a) = -\gamma'(b)$. In particular, there does not exist an angle φ such that both T_φ and $T_{\varphi+\pi}$ contain an inner point of γ. Furthermore, there exists an angle θ such that $\gamma \cap T_\theta = \{\gamma(0)\}$ and $\gamma \cap T_{\theta+\pi} = \{\gamma(L)\}$.*

Proof. Assume that there is an angle φ such that $\gamma(a) \in T_\varphi$, $\gamma(b) \in T_{\varphi+\pi}$, and $\gamma'(a) = -\gamma'(b)$ and assume without loss of generality that $a < b$. Let $\widehat{\mathbf{m}}(x) = (-\sin x, \cos x)$. If $0 < a$, we must have $\gamma'(a) = \widehat{\mathbf{m}}(\varphi)$ and hence $\gamma'(b) = \widehat{\mathbf{m}}(\varphi + \pi)$. The same holds when $b < L$. If $a = 0$ and $b = L$, we must have $\gamma'(0) \in \widehat{\mathbf{m}}([\varphi, \varphi + \pi))$ since otherwise the interior of $\overline{\gamma}$ is on wrong side of T_φ in a neighborhood around $\gamma(0)$. Similarly $\gamma'(L) \in \widehat{\mathbf{m}}((\varphi, \varphi + \pi])$. Thus, we must have $\gamma'(0) = \widehat{\mathbf{m}}(\varphi)$ and $\gamma'(L) = \widehat{\mathbf{m}}(\varphi + \pi)$. We have seen that in any case, $\gamma'(a) = \widehat{\mathbf{m}}(\varphi)$ and $\gamma'(b) = \widehat{\mathbf{m}}(\varphi + \pi)$.

Construct the curve ρ consisting of $\gamma([a, b])$, then a line segment ℓ_0 on $T_{\varphi+\pi}$ with direction $\gamma'(b)$ of length L, then a counterclockwise semi-circle $\mathcal{C}$ tangential to T_φ and $T_{\varphi+\pi}$, and then a line segment ℓ_1 on T_φ to $\gamma(a)$. We now observe that ρ is simple. $\gamma([a, b])$ cannot cross ℓ_0 or ℓ_1 due to Lemma 8 and not $\mathcal{C}$ since ℓ_0 has length L. Furthermore, ρ is continuously differentiable.

Let ϑ_ρ and L_ρ be the tangent indicatrix associated to and length of ρ, respectively. Hopf's Umlaufsatz gives that $\vartheta_\rho(L_\rho) - \vartheta_\rho(0) = 2\pi$, see for instance [11, Thm. 11.1]. We now have $\vartheta_\rho(L_\rho) - \vartheta_\rho(b - a) = \pi$ since ϑ_ρ increases by π along $\mathcal{C}$. Hence $\vartheta_\gamma(b) - \vartheta_\gamma(a) = \vartheta_\rho(b - a) - \vartheta_\rho(0) = 2\pi - (\vartheta_\rho(L_\rho) - \vartheta_\rho(b - a)) = \pi$. By Lemma 5, either $b - a \geq \pi$ or γ is not in $\mathcal{K}^+$, which is a contradiction.

If T_φ contains an inner point $\gamma(t)$, then $\gamma'(t) = \widehat{\mathbf{m}}(\varphi)$. Therefore, $T_{\varphi+\pi}$ cannot contain an inner point, i.e., $\gamma \cap T_{\varphi+\pi} \subset \{\gamma(0), \gamma(L)\}$. Hence, there exists an angle α such that $\gamma \cap T_\alpha \subset \{\gamma(0), \gamma(L)\}$. Assume without loss of generality that $\gamma(0) \in \gamma \cap T_\alpha$. Now, let β be the smallest angle larger than α such that T_β contains an inner point of γ or $\widehat{\mathbf{m}}(\beta) = \gamma'(0)$. Then $\gamma(0) \in \gamma \cap T_\beta$ and $\gamma \cap T_{\beta+\pi} = \{\gamma(L)\}$. Thus, by choosing θ slightly smaller than β, where θ depends on γ, we also have $\gamma \cap T_{\theta+\pi} = \{\gamma(L)\}$ while $\gamma \cap T_\theta = \{\gamma(0)\}$. ◀

▶ **Lemma 10.** *If $L < \pi$, then $\partial D[\gamma(0), \gamma(L)]$ spans an angle of less than π, $\delta > 1$, and γ is contained in $\overline{U}$, where U is the open unit disk with a center in the closed half-plane to the left of the directed line from $\gamma(0)$ through $\gamma(L)$ such that $\gamma(0), \gamma(L) \in \partial U$.*

Proof. If the interior arc of ∂D spans an angle of π or more, it contains two antipodal points defining a diameter of D. Assume without loss of generality that the diameter is horizontal. Then, the tangents T_0 and T_π contain points on γ with opposite vertical tangents, contradicting Lemma 9.

To prove the two remaining statements, we show that we can reduce the problem to the case where $\delta = 1$. Note that in that case, $D = U$. Assume for contradiction that one of the remaining statements does not hold. In each case, γ is not contained in $\overline{U}$. We choose a maximal interval $[p, q] \subseteq [0, L]$ such that $\gamma([p, q]) \cap \overline{U} = \{\gamma(p), \gamma(q)\}$. Considering U instead of D and restricting γ to the interval $\gamma([p, q])$, we now have a similar situation as described in the Lemma where $\delta = 1$ and $\gamma([p, q])$ is not contained in U. Hence, let us in the following assume that $\delta = 1$ and that $\gamma([0, L])$ is not a subset of $\overline{D}$.

Denote $\mathcal{I} = \partial D[\gamma(0), \gamma(L)]$. Let θ be an angle such that $\gamma \cap T_\theta = \{\gamma(0)\}$ and $\gamma \cap T_{\theta+\pi} = \{\gamma(L)\}$. The existence of θ is ensured by Lemma 9.

Now, imagine that we slide $\mathcal{I}$ along the two tangents T_θ and $T_{\theta+\pi}$ in direction $\widehat{\mathbf{m}}(\theta)$ until the last point of contact with R. Let $\mathcal{I}'$ be the translation of $\mathcal{I}$ to this point and $\gamma(\sigma)$ the last point of contact, which are well-defined since both γ and $\mathcal{I}$ are compact. $\gamma(\sigma)$ cannot be any of the end points of $\mathcal{I}'$ since the tangents only contain end points of γ. Hence, $\gamma(\sigma)$ is an inner point of $\mathcal{I}'$, and $\mathcal{I}'$ and γ share a tangent in $\gamma(\sigma)$. It follows that $\mathcal{I}'$ is an arc on the boundary of the open unit disk A with center $\gamma(\sigma) + \mathbf{n}(\sigma)$. γ enters A on a point on the arc from $\gamma(\sigma)$ to the end of $\mathcal{I}'$ since the end of $\mathcal{I}'$ is not on γ. That contradicts Lemma 6. ◀

▶ **Theorem 11.** *The region R contains an open unit disk U such that $\gamma(s) \in \partial U$ for some $s \in [\sigma, L-\sigma]$, where $\sigma = \min\left\{\frac{\|\gamma(0)-\gamma(L)\|}{2}, \frac{\pi}{2}\right\}$.*

Proof. Inspired by Ahn et al. [1], we proceed by induction on $\lfloor L/\pi \rfloor$. For the base-case, assume that $L < \pi$.

We see that for any $s \in (0, L)$, we can choose an open unit disk $U \subset R$ such that $\gamma(s) \in \partial U$. Namely, assume that for some $p \in (0, L)$ there is no such unit disk. Consider the largest open unit disk $D' \subset R$ where $\gamma(p) \in \partial D'$. From Lemma 7, we know that $\partial D'$ contains a point $\gamma(q) \in \gamma([0, L]) \setminus \{\gamma(p)\}$. Lemma 6 gives that γ is disjoint from $\overline{D'}$ in a neighborhood around p since the radius of D' is less than 1. Therefore, we assume without loss of generality that $p < q$ and that $\gamma([p, q]) \cap \overline{D'} = \{\gamma(p), \gamma(q)\}$. But then $\gamma[p, q]$ is not in $\mathcal{K}^+$ by Lemma 10.

For the induction step, assume that the theorem is true when $\lceil L/\pi \rceil \leq k-1$ and consider a curve γ with length L such that $\lceil L/\pi \rceil = k$. We first show the weaker statement that there is an open unit disk $U \subset R$ such that $\gamma(s) \in \partial U$ for some $s \notin \{0, L\}$. This will be needed when showing that we can choose $s \in [\sigma, L-\sigma]$. Let $p = L/2$ and let D' be the largest open disk such that $D' \subset R$ and $\gamma(p) \in \partial D'$. If the radius δ' of $\partial D'$ is at least 1, we are done. Otherwise, due to Lemma 7, $\partial D'$ contains a point $\gamma(q) \in \gamma([0, L]) \setminus \{\gamma(p)\}$.

Lemma 6 gives that γ is disjoint from $\overline{D'}$ in an interval to each side of p since $\delta' < 1$. Therefore, we assume without loss of generality that $p < q$ and that $\gamma((p, q)) \cap \overline{D'} = \emptyset$. The curve $\gamma([p, q])$ has length at most $L/2$, so the induction hypothesis gives that the open region R' bounded by $\partial D'[\gamma(q), \gamma(p)]$ and $\gamma([p, q])$ contains an open unit disk U, such that ∂U contains $\gamma(s)$ from some $s \in (p, q)$. We note that $R' \subseteq R$ and $s \notin \{0, L\}$, as required.

To prove that we can choose s in the interval $[\sigma, L-\sigma]$, let D_1 be the largest open disk contained in R such that ∂D_1 contains $\gamma(\sigma)$. Assume D_1 has a radius smaller than 1, since otherwise, we are done. By Lemma 7, there exists a point $\gamma(p_1) \in \partial D_1 \setminus \{\gamma(\sigma)\}$. By Lemma 10 we get that $p_1 > \sigma$ since otherwise $\gamma([0, \sigma])$ is not in $\mathcal{K}^+$. If $p_1 \leq L - \sigma$, we know from the above discussion that there exists an open unit disk U contained in the region bounded by $\partial D_1[\gamma(p_1), \gamma(\sigma)]$ and $\gamma([\sigma, p_1])$ such that ∂U contains a point $\gamma(s) \in \gamma((\sigma, p_1))$. Hence also $s \in [\sigma, L-\sigma]$, as stated in the theorem.

If $p_1 > L - \sigma$, let R_1 be the region bounded by $D_1[\gamma(p_1), \gamma(\sigma)]$ and $\gamma([\sigma, p_1])$, and let D_2 be the largest open disk contained in R_1 such that ∂D_2 contains $\gamma(L-\sigma)$. Assume that D_2 has a radius of less than 1, since otherwise we are done. Lemma 7 gives that there is a point $\gamma(p_2) \in \partial D_2 \cap \gamma([\sigma, p_1]) \setminus \{\gamma(L-\sigma)\}$. Lemma 10 gives that $p_2 < L - \sigma$ since otherwise $\gamma([p_2, p_1])$ is not in $\mathcal{K}^+$. Hence, $p_2 \in [\sigma, L-\sigma)$, and from the above discussion we know that there is an open unit disk U contained in the open region bounded by $\partial D_2[\gamma(L-\sigma), \gamma(p_2)]$ and $\gamma([p_2, L-\sigma])$ such that ∂U contains a point $\gamma(s) \in \gamma((p_2, L-\sigma))$. This proves the theorem. ◀

2.1.1 A generalization to curves with sharp corners

Let $\gamma : [0, L] \to \mathbb{R}^2$ be a curve and assume that the backward tangent $\gamma'_-(s) = \lim_{x \to s-} \frac{\gamma(s)-\gamma(x)}{s-x}$ is defined for all $s \in (0, L]$ and the forward tangent $\gamma'_+(s) = \lim_{x \to s+} \frac{\gamma(s)-\gamma(x)}{s-x}$ is defined for all $s \in [0, L)$. Furthermore, assume that $\|\gamma'_-(s)\|$ and $\|\gamma'_+(s)\|$ is 1 where they are defined.

We say that $\gamma(x)$ is a *corner* if $\gamma'_+(x) \neq \gamma'_-(x)$. For technical reasons, we also define $\gamma(0)$ and $\gamma(L)$ to be corners and say that $\gamma(x)$ is a *proper* corner if $x \notin \{0, L\}$. Let X be the set of all corners. We do not assume X to be finite, but it follows from the Denjoy–Young–Saks Theorem [4, p. 63, Thm. 4.2] that X is countable.

We define $\mathbf{n}_+(x) = \widehat{\gamma'_+}(x)$ for every $x \in [0, L)$ and $\mathbf{n}_-(x) = \widehat{\gamma'_-}(x)$ for every $x \in (0, L]$. At a point $\gamma(x)$ which is not a corner, the unit-normal $\mathbf{n}(x) = \mathbf{n}_+(x) = \mathbf{n}_-(x)$ is defined as for differentiable curves. At other points, the unit-normal is undefined.

At a proper corner $\gamma(x)$, γ either makes a turn to the left or to the right. If $\langle \widehat{\gamma'_-}(x), \gamma'_+(x) \rangle < 0$, it is a left-turn and if $\langle \widehat{\gamma'_-}(x), \gamma'_+(x) \rangle > 0$ it is a right-turn. If $\langle \widehat{\gamma'_-}(x), \gamma'_+(x) \rangle = 0$, we have $\gamma'_-(x) = -\gamma'_+(x)$. To decide to which side γ turns in this case, we consider for each $n \in \mathbb{N}$ a circle $\mathcal{C}_n$ with center $\gamma(x)$ and a radius so small that it intersects both $\gamma([x - 1/n, x])$ in a point y_n and $\gamma([x, x + 1/n])$ in a point z_n. Since $\gamma'_-(s)$ and $\gamma'_+(s)$ are well defined, the size of the angle of the arc on $\mathcal{C}_n$ from y_n counterclockwise to z_n converges to either 0 or 2π, in which case γ makes a right- or a left-turn at $\gamma(x)$, respectively.

For each corner $\gamma(x) \in X$ where $x < L$, we define $L_x = (x, \inf_{x<y,\gamma(y)\in X} y)$, where L_x is possibly empty. Note that γ is differentiable on each interval L_x. We say that γ has *bounded left-curvature* and write $\gamma \in \mathcal{L}^+$ if each proper corner is a right-turn and $\gamma(\overline{L_x}) \in \mathcal{K}^+$ for every interval L_x. Since X is countable, γ' is defined and continuous almost everywhere. Therefore $\|\gamma'(y)\| = 1$ almost everywhere. Hence $\int_0^\ell \|\gamma'(y)\| dy = \ell$, so it makes sense to say that γ is parameterized by arc length. It also holds that the shortest curve in $\mathcal{L}^+$ between two points follows the line segment between the points, for instance, the proof outlined in [3, Exercise 1.2] generalizes to our case. The class $\mathcal{K}^+$ is clearly a proper subset of $\mathcal{L}^+$.

We are now ready to prove Theorem 1 as stated in Section 1.

Proof of Theorem 1. We prove that Theorem 11 also holds when $\gamma \in \mathcal{L}^+$ instead of $\gamma \in \mathcal{K}^+$. We construct a sequence of curves $\gamma_1, \gamma_2, \ldots$ in $\mathcal{K}^+$ with lengths $L_1, L_2, \ldots$, such that $\gamma_n(0) = \gamma(0)$ and $\gamma_n(L_n) = \gamma(L)$ for all n. The sequence converges to γ in such a way that Theorem 11 applied to each curve γ_n implies the existence of a unit disk in R as wanted. Informally, we may describe γ_n as the result of rounding the proper corners of γ with circular arcs of radii less than $1/2n$.

Each curve γ_n is itself the limit of a sequence $\gamma_{n,0}, \gamma_{n,1}, \ldots$ of curves in $\mathcal{L}^+$, such that $\gamma_{n,0} = \gamma$ and $\gamma_{n,m+1}$ is the result of rounding at least one corner of $\gamma_{n,m}$. The endpoints of $\gamma_{n,m}$ are $\gamma(0)$ and $\gamma(L)$.

We define $\gamma_{n,0} = \gamma$ and $C_0 = \emptyset$. C_m consists of the circular arcs on $\gamma_{n,m}$ we have added to round the corners of γ. Let $x_1, x_2, \ldots$ be an enumeration of the parameters such that $\gamma(x_1), \gamma(x_2), \ldots$ are all the proper corners of γ. Assume that we have already defined the curve $\gamma_{n,m}$. We now show how $\gamma_{n,m+1}$ is defined.

Choose the smallest k such that $\gamma(x_k)$ is also a corner of $\gamma_{n,m}$ and define x such that $\gamma(x_k) = \gamma_{n,m}(x)$. Let W we a wedge with tip $\gamma_{n,m}(x)$ such that $\gamma_{n,m}$ is contained in W in an interval before and after x. Such a wedge exists since the forward and backward tangents of $\gamma_{n,m}$ are well-defined in x. Let $w > 0$ be a number so small that the intersection between $\gamma_{n,m}$ and any open disk with center $\gamma_{n,m}(x)$ and radius at most w is an interval of $\gamma_{n,m}$ contained in W. Then, let

$$r = \min\left\{1/n, w, \operatorname{dist}(\gamma_{n,m}(x), C_m), \|\gamma_{n,m}(x) - \gamma(0)\|, \|\gamma_{n,m}(x) - \gamma(L)\|\right\}.$$

We assume inductively that $\operatorname{dist}(\gamma_{n,m}(x), C_m) > 0$ and hence $r > 0$. Let $\mathcal{A}$ be the open disk with center x and radius r, and let $\gamma_{n,m}([a, b])$ be the interval $\overline{\mathcal{A}} \cap \gamma_{n,m}$, where clearly $x \in [a, b]$. The interval $\gamma_{n,m}([a, x])$ divides $W \cap \overline{\mathcal{A}}$ into two pieces, and since $\gamma_{n,m}$ makes a right-turn at $\gamma_{n,m}(x)$, $\gamma_{n,m}([x, b])$ is contained in the piece to the right of $\gamma_{n,m}([a, x])$.

Let $\mathcal{B}$ be the open disk with center $\gamma_{n,m}(x)$ and radius $r/2$ and let $\partial\mathcal{B}'$ be the arc $\partial\mathcal{B} \cap W$. The intersection points between $\partial\mathcal{B}'$ and $\gamma_{n,m}$ divide $\partial\mathcal{B}'$ into smaller arcs. Consider the last intersection point $\gamma_{n,m}(c)$ between $\gamma_{n,m}([a, x])$ and $\partial\mathcal{B}'$ following $\partial\mathcal{B}'$ in counterclockwise

direction. The arc $\partial\mathcal{B}''$ on $\partial\mathcal{B}'$ beginning at $\gamma_{n,m}(c)$ must end at a point $\gamma_{n,m}(d)$ on $\gamma_{n,m}([x,b])$, since otherwise, $\gamma_{n,m}$ cannot leave $\mathcal{B}$ after $\gamma_{n,m}(x)$.

For each point $p \in \partial\mathcal{B}''$, we consider the largest open disk $\mathcal{C}_p$ with center p disjoint from $\gamma_{n,m}$. In the following, we show that for each $p \in \partial\mathcal{B}''$, we have $\mathcal{C}_p \subset \mathcal{A}$ and that $\gamma_{n,m}(x) \notin \partial\mathcal{C}_p$: Let $\mathcal{D}$ be the largest open disk with center p contained in W and let S be the region bounded by $\gamma_{n,m}([a,b])$ and the arc on $\partial\mathcal{A}$ from $\gamma_{n,m}(a)$ counterclockwise to $\gamma_{n,m}(b)$. Then $\mathcal{C}_p \subset S \subset W$, so $\mathcal{C}_p \subset \mathcal{D}$. But $\partial\mathcal{D}$ does not contain $\gamma_{n,m}(x)$ since $\gamma_{n,m}(x)$ is the tip of the wedge W. Hence the radius of $\mathcal{C}_p$ is less than $r/2$ for all $p \in \partial\mathcal{B}''$.

When p is close to $\gamma_{n,m}(c)$, $\partial\mathcal{C}_p$ contains a point on $\gamma_{n,m}([a,x))$ while when p is close to $\gamma_{n,m}(d)$, $\partial\mathcal{C}_p$ contains a point on $\gamma((x,b])$. Therefore, there must be a point $P \in \partial\mathcal{B}''$ such that $\mathcal{C}_P$ contains a point $\gamma_{n,m}(e) \in \gamma_{n,m}([a,x))$ and a point on $\gamma_{n,m}(f) \in \gamma_{n,m}((x,b])$.

The interior of S, and hence $\mathcal{C}_P$, is to the right of $\gamma_{n,m}([a,b])$ since $\gamma_{n,m}(x)$ is a right-turn. Therefore, since all corners on $\gamma_{n,m}$ are right-turns, $\gamma_{n,m}(e)$ and $\gamma_{n,m}(f)$ cannot be corners, so $\gamma_{n,m}$ is differentiable in e and f, and $\partial\mathcal{C}_P$ and $\gamma_{n,m}$ have the same tangents there. We now define $\gamma_{n,m+1}$ as the curve identical to $\gamma_{n,m}$, but where $\gamma_{n,m+1}$ follows the arc $\partial\mathcal{C}'_P$ on $\partial\mathcal{C}_P$ from $\gamma_{n,m}(e)$ clockwise to $\gamma_{n,m}(f)$ instead of the interval $\gamma_{n,m}([e,f])$. We note that $\partial\mathcal{C}'_P$ is disjoint from all the arcs in C_m and define $C_{m+1} = C_m \cup \partial\mathcal{C}'_P$. We also see that $\operatorname{dist}(\gamma_{n,m+1}(y), C_{m+1}) > 0$ for every remaining corner $\gamma_{n,m+1}(y)$ on $\gamma_{n,m+1}$.

We now define γ_n as $\bigcap_{m=0}^{\infty} \gamma_{n,m}$, i.e., γ_n consists of all the points on γ which are never substituted by an arc and all the arcs $\bigcup_{m=1}^{\infty} C_m$. It now follows from the construction that $\gamma_n \in \mathcal{K}^+$, since each point on γ_n is either a point on a clockwise circular arc or a point on γ which is not a corner.

We note that the Hausdorff distance between γ and γ_n converges to 0, i.e., that the maximum distance from a point on γ to the closest point on γ_n converges to 0 as $n \to \infty$ and likewise for the maximum distance from a point on γ_n to the closest point on γ.

By Theorem 11, there exists a unit disk U_n in the region R_n bounded by D and γ_n for each n, such that U_n touches γ_n in an inner point $\gamma_n(s_n)$, where $s_n \in [\sigma, L_n - \sigma]$, where $\sigma = \min\left\{\frac{\|\gamma(0)-\gamma(L)\|}{2}, \frac{\pi}{2}\right\}$. Hence, there exists an accumulation point c among the sequence of the centers of the disks $U_1, U_2, \ldots$. Let U be the open unit disk with center c. Since the Hausdorff distance between γ and γ_n converges to 0, U is contained in R. It remains to be proven that $\gamma(s) \in \partial U$ for some $s \in [\sigma, L - \sigma]$.

We call a point on γ_n a *fixed point* if the point also exists on γ. We note that the distance along γ_n from an arbitrary point on γ_n to the closest fixed point in each direction is less than π/n, since the point which are not fixed appear on circular arcs of lengths less than π/n. Furthermore, we see that every arc $\mathcal{C}'_P$ is not longer than the interval $\gamma_{n,m}([e,f])$ that it substitutes: It is well-known that $\gamma([e,f])$ is at least as long as any convex polygonal chain from e to f contained in the area bounded by $\gamma[(e,f)]$ and the line segment from $\gamma(e)$ to $\gamma(f)$ [15]. The arc $\mathcal{C}'_P$ is the limit of a converging sequence of such polygonal chains and is hence also not longer than $\gamma_{n,m}([e,f])$. Therefore, for each fixed point $\gamma_n(v) = \gamma(w)$, we have $v < w$ and $L_n - v < L - w$.

Let (a_n) be a subsequence of the positive integers (n) such that the center of U_{a_n} converges to c. The sequence $(\gamma_{a_n}(s_{a_n}))$ has an accumulation point $\gamma(s)$. Let (b_n) be a subsequence of (a_n) such that $\gamma_{b_n}(s_{b_n})$ converges to $\gamma(s)$. The sequence (s_{b_n}) has an accumulation point χ. Let (c_n) be a subsequence of (b_n) such that s_{c_n} converges to χ. We choose a sequence (χ_n) such that (i) $\chi_n \le s_{c_n}$, (ii) χ_n converges to χ, and (iii) $(\gamma_{c_n}(\chi_n))$ is a sequence of fixed points. As mentioned above, we have $\chi_n \le s$ for each n since $\gamma(\chi_n)$ is a fixed point. We thus have

$$\sigma \le \lim_{n\to\infty} s_{c_n} = \chi = \lim_{n\to\infty} \chi_n \le s.$$

Similarly, the sequence $(L_n - s_{b_n})$ has an accumulation point ψ, and we choose a subsequence (d_n) of (b_n) such that $L_n - s_{d_n}$ converges to ψ and a sequence (ψ_n) such that (i) $\psi_n \geq s_{d_n}$, (ii) $L_n - \psi_n$ converges to ψ, and (iii) $(\gamma_{d_n}(\psi_n))$ is a sequence of fixed points. We then have

$$\sigma \leq \lim_{n\to\infty} L_n - s_{d_n} = \psi = \lim_{n\to\infty} L_n - \psi_n \leq L - s.$$

◀

2.1.2 Closed curves with bounded convex curvature

We say that a closed, simple curve γ has *bounded convex curvature* if every closed, connected, proper subset of γ can be parameterized as a curve $\tilde{\gamma}$ in $\mathcal{L}^+$ following the counterclockwise direction of γ.

Proof of Corollary 2. Consider a maximal open disk D in the interior of γ in the sense that no other open disk in the interior of γ contains D. Assume that D has radius less than 1. Clearly, ∂D contains at least two points $a, b \in \gamma$. Since D has radius less than 1, γ must be different from ∂D. Therefore, we can choose a and b such that $\gamma[a, b] \cap \overline{D} = \{a, b\}$. Now, Theorem 1 implies that the curve defined by $\partial D[b, a]$ and $\gamma[a, b]$ contains an open unit-disk in its interior. ◀

3 Algorithm

3.1 Preliminaries

We assume that a simply-connected region P_1 in the plane is given such that the boundary ∂P_1 consists of a finite number of line segments and circular arcs of arbitrary radii. For ease of presentation, we assume that ∂P_1 is a simple curvilinear polygon, but our algorithm works as long as ∂P_1 is weakly simple. We can therefore without loss of generality assume that P_1 is an open set. We set $P \leftarrow P_1$ and our algorithm keeps removing parts of P while maintaining the invariant that $\overline{P}$ contains every subset of $\overline{P_1}$ of bounded convex curvature. In the end, $\overline{P}$ itself has bounded convex curvature and it follows that $\overline{P}$ is the unique maximum subset of $\overline{P_1}$ of bounded convex curvature.

The region P is always a collection of disjoint, open, simply-connected sets each of which is bounded by a simple curvilinear polygon. P is represented by its boundary ∂P, which is a collection of disjoint, closed, simple curves where no curve is contained in the interior of another. The input P_1 is defined by one such curve. The open region enclosed by each curve is one connected component of P. ∂P is represented as a set of points known as the *vertices* of ∂P and a set of line segments and circular arcs known as the *arcs* of ∂P. We think of line segments as circular arcs with infinite radius and therefore in most cases use the word *arcs* for both circular arcs and line segments. Depending on the context, we may consider a vertex as a point or a set containing a single point. An *object* of ∂P is a vertex or an arc. We use the convention that an arc includes its endpoints. Every two arcs of ∂P are disjoint except possibly at the endpoints, and for each vertex there are two arcs having an endpoint at that vertex. This way, the arcs form the closed curves bounding P. We always use n to denote the number of vertices of the input ∂P_1.

The boundary of each connected component of P is oriented in counterclockwise direction and we say that a point moving on ∂P is moving *forward* (resp. *backward*) if it is following (resp. not following) the orientation of the boundary. Similarly, we orient every arc following the orientation of the boundary of the component containing it. We denote the endpoints

of an arc A as $s(A)$ and $t(A)$, so that a point moving forward along A moves from $s(A)$ to $t(A)$. An arc which turns to the left when traversed in forward direction is called a *convex arc*. A *concave arc* is defined analogously. Line segments are regarded as both concave and convex arcs at the same time. A vertex v is *convex* if the interior angle of ∂P at v is strictly less than π. If the angle is strictly more than π, v is *concave*.

Let A be an arc of ∂P and $a \in A$ a point on A. Then $\mathbf{n}_A(a)$ is the unit-normal of A at a which points to the left relative to the orientation of A. We say that two arcs A and B are *tangential* if $t(A) = s(B)$ and $\mathbf{n}_A(t(A)) = \mathbf{n}_B(s(B))$. Note that A and B are tangential if and only if the vertex $t(A)$ is neither convex nor concave.

3.2 High-level description of the algorithm

A high-level description of the algorithm is given as Algorithm 1. The basic format is to maintain a stack Σ of pointers to objects of ∂P causing the convex curvature condition on ∂P to be violated. An object $\sigma \in \Sigma$ can be a convex vertex, a convex arc of radius less than 1, or a special *cut arc* of radius 1, where one or both of the endpoints might be convex vertices. In each iteration of the loop at line 3, we test if the object σ on the top of Σ is still on ∂P (it might have been removed in another iteration) and if so, we eliminate it by removing from P a subset $\mathcal{V} \subseteq P$. The object σ appears on the boundary $\partial\mathcal{V}$ and not on the boundary of $P \setminus \mathcal{V}$. By *performing a cut* or simply a *cut*, we mean the process of removing $\mathcal{V}$ from P. It is important to choose $\mathcal{V}$ such that $\mathcal{V} \cap Q = \emptyset$ for every set $Q \subset \overline{P_1}$ of bounded convex curvature. Theorem 1 will be used to prove this.

It is possible that a cut splits P into more components. The algorithm will then keep working on each component separately. A cut can introduce new unit-radius cut arcs on ∂P, the endpoints of which can be convex vertices. Therefore, these new cut arcs are added to Σ so that the new convex vertices are eventually removed. The algorithm terminates when Σ is empty, which means that P has bounded convex curvature.

Algorithm 1: `SubsetOfBoundedConvexCurvature`(P)

```
1 Add all convex arcs of ∂P with a radius less than 1 to Σ.
2 Add all convex vertices of ∂P to Σ.
3 while Σ ≠ ∅
4 |   Let σ be the topmost element on Σ and remove σ from Σ.
5 |   if σ ⊂ ∂P and σ is not a perfect cut arc
6 |   |   Cut away V from P and let C_j, j = 1,...,t, be the new cut arcs after the cut.
7 |   |   Add each new cut arc C_j to Σ.
8 return P
```

Let P_i be the set P in the beginning of iteration $i = 1, 2, \ldots$ of the loop at line 3 in Algorithm 1. We shall prove that the algorithm terminates after $k = O(n)$ iterations. Hence we have

$$P_1 \supseteq P_2 \supseteq \cdots \supseteq P_k.$$

The challenge is to define $\mathcal{V}$ in each iteration so that $k = O(n)$, $P \leftarrow P \setminus \mathcal{V}$ can be computed efficiently, and $\mathcal{V} \cap Q = \emptyset$ for every set $Q \subset \overline{P_1}$ of bounded convex curvature.

3.3 Specifying a cut using an arc $\mathcal{C}$

Let σ be an object from Σ that is to be eliminated in a certain iteration of Algorithm 1, and let K be the connected component of P such that $\sigma \subset \partial K$. If we conclude that K does not contain any non-empty set of bounded convex curvature, we set $\mathcal{V} = K$, so that all of the component K is removed from P. Otherwise, we specify $\mathcal{V}$ using a circular arc $\mathcal{C}$ of unit radius. See Figure 7, where $\mathcal{C}$ is grey and $\mathcal{V}$ is the area enclosed by the thick closed curve. Assume for now that we have defined $\mathcal{C}$. Let $s(\mathcal{C})$ and $t(\mathcal{C})$ be the points where $\mathcal{C}$ starts and ends in counterclockwise direction, respectively. The endpoints are included in $\mathcal{C}$ and are points on ∂K. They will be defined so that $\sigma \subseteq \partial K[s(\mathcal{C}), t(\mathcal{C})]$. It will follow from the definition of $\mathcal{C}$ that there will be at most two vertices on $\partial K(s(\mathcal{C}), t(\mathcal{C}))$. For efficiency, the algorithm may choose an arc $\mathcal{C}$ that intersects $\partial K(s(\mathcal{C}), t(\mathcal{C}))$.

The arc $\mathcal{C}$ divides K into open regions $R_1, \ldots, R_r$, which are the connected components of $K \setminus \mathcal{C}$. Assume without loss of generality that among all these regions, exactly the regions $R_1, \ldots, R_s$, $s \leq r$, contain a part of $\partial K(s(\mathcal{C}), t(\mathcal{C}))$ on their boundary. We note that $s = 1$ if $\mathcal{C}$ does not intersect $\partial K(s(\mathcal{C}), t(\mathcal{C}))$, as is the case in Figure 7. Define $\mathcal{V}$ as the union $\bigcup_{i=1}^{s} K \cap \overline{R_i}$. Then $K \setminus \mathcal{V}$ is open and hence so is our new $P \leftarrow P \setminus \mathcal{V}$. The arc $\mathcal{C}$ will be carefully chosen so that $\mathcal{V} \cap Q = \emptyset$ for every set $Q \subset \overline{P_1}$ of bounded convex curvature.

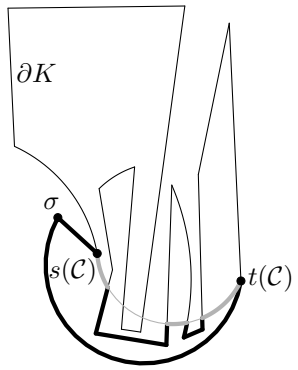

Figure 7

▶ **Lemma 12.** *$\mathcal{V}$ is connected and for each $i = 1, 2, \ldots$, the boundary ∂K of each connected component K of P_i is a simple curvilinear polygon.*

Proof. Let two points $p, q \in \mathcal{V}$ be given. We can connect p to q by a curve contained in $\mathcal{V}$ by walking along $\partial K(s(\mathcal{C}), t(\mathcal{C}))$. Hence, $\mathcal{V}$ is connected.

Assume that every component of P_i has a boundary which is a simple curve and assume that the same is not true for the connected component K of P_{i+1}. Let $\mathcal{C}$ be the arc specifying the cut that be made in P_i to obtain P_{i+1}. There must be an arc $\mathcal{C}[a, b]$ on ∂K and a point $x \in \mathcal{C}(a, b)$ such that $x \in \partial K(b, a)$. But that implies that K is not connected, which is a contradiction. ◀

Consider the case where we remove a proper subset $\mathcal{V}$ from a connected component K of P_i to obtain P_{i+1}. One or more arcs $\mathcal{C}_1, \ldots, \mathcal{C}_t$ of ∂P_{i+1} are subsets of $\mathcal{C}$ and not arcs of ∂P_i. We denote the arcs $\mathcal{C}_1, \ldots, \mathcal{C}_t$ as *new cut arcs* of ∂P_{i+1}. An arc of ∂P_{i+1} which is a subset of a new cut arc of ∂P_j, $j \leq i$, is not a new cut arc of ∂P_{i+1}. An arc of ∂P_{i+1} is a *cut arc* if it is the subset of a new cut arc of ∂P_j for some $j \leq i + 1$. We say that a cut arc $\mathcal{C}$ of ∂P is *perfect* if none of the endpoints of $\mathcal{C}$ are convex vertices.

▶ **Lemma 13.** *Every cut arc of ∂P is convex.*

Proof. We show that if the statement is true for P_i, then it is also true for P_{i+1}. Let $\mathcal{C}$ be the arc specifying a cut to be performed in P_i. Assume that $\mathcal{V}$ is a subset of the connected component K in P_i. We proceed by induction on the number r of connected components $R_1, \ldots, R_r$ of $K \setminus \mathcal{C}$. If $r = 1$, the statement is trivial, since $\mathcal{V} = R_1 = K$.

Consider the case $r = 2$. Assume that $\mathcal{V} = K \cap \overline{R}_1$ is removed from K and that R_2 remains, since otherwise there is nothing to show. R_2 is bounded by some interval $\partial K[a, b]$ and a subset $\mathcal{C}'$ of $\mathcal{C}$ with endpoints a and b. The points must appear in the order $s(\mathcal{C}), t(\mathcal{C}), a, b$ along ∂K, since otherwise R_2 is removed. Assume for contradiction that $\mathcal{C}'$ is a clockwise arc on ∂R_2. That means that the points appear in the order $s(\mathcal{C}), a, b, t(\mathcal{C})$ along $\mathcal{C}$. Furthermore,

since $r = 1$, $\mathcal{C}[s(\mathcal{C}), a]$ and $\mathcal{C}[b, t(\mathcal{C})]$ are disjoint from K, which means that $\partial K[b, a]$ does not cross $\mathcal{C}$ from one side to the other. There exists a point $a' \in \partial K[t(\mathcal{C}), a]$ which is not on $\mathcal{C}$ since otherwise $b \in \partial K[t(\mathcal{C}), a]$, contradicting that ∂K is simple. Let $s'(\mathcal{C})$ be the first point on $\mathcal{C}[s(\mathcal{C}), a]$ we meet when traveling along ∂K from b. Let b' be the last point on $\mathcal{C}[a, t(\mathcal{C})]$ we meet when traveling along ∂K from b to $s'(\mathcal{C})$. Clearly, $b' \neq a$ and $b' \neq t(\mathcal{C})$ since ∂K is simple. By construction, $\partial K[b', s'(\mathcal{C})]$ and $\mathcal{C}[s'(\mathcal{C}), b']$ form a simple closed curve J. Orient J so that $\mathcal{C}[s'(\mathcal{C}), b']$ is a counterclockwise arc when traversing J in forward direction. When walking from a to a' along $\partial P[a', a]$, we leave J to its left side. When walking from b' to $t(\mathcal{C})$, we leave J to its right side. Hence J separates $t(\mathcal{C})$ from a'. Therefore, $\partial K[t(\mathcal{C}), a']$ crosses J, which is a contradiction.

Assume now $r \geq 3$ and that the statement is true for smaller values of r. Since P is simply connected, the regions $R_1, \ldots, R_r$ induce a tree T such that two regions R_j and R_k are connected if there is an arc $\mathcal{C}(x, y) \subset \partial R_j \cap \partial R_k$ for $x \neq y$. Let R_j and R_k be two leafs in T. Let $K_j = K \setminus \overline{R_j}$ and $K_k = K \setminus \overline{R_k}$. K_j and K_k are both subdivided into $r - 1$ regions by $\mathcal{C}$, and let K'_j and K'_k be the result of performing the cut specified by $\mathcal{C}$ in K_j and K_k, respectively. Note that R_j and R_k are not neighbors in T since $r \geq 3$ and hence both of the unique neighbors in T of R_j and R_k are also regions in both $K_j \setminus \mathcal{C}$ and $K_k \setminus \mathcal{C}$. Therefore, every new cut arc in P_{i+1} is a new cut arc in either K'_j or K'_k. Hence, the induction hypothesis gives that every cut arc of ∂P_{i+1} is convex. ◀

3.4 The bisector curve ξ

A specific type of bisector curve, illustrated in Figure 8, turns out to be useful when describing how to define the arc $\mathcal{C}$ specifying a cut and when proving the correctness of our algorithm. The curve is somewhat related to the medial axis of ∂P. Before introducing the curve, we need the concept of an *osculating circle*.

Let A be an arc of ∂P. A circle C *osculates* A at $a \in A$ if the interior of C is disjoint from A, C and A have the point a in common, and the center x of C is a point $a + \mathbf{n}_A(a) \cdot c$ where $c \geq 0$. Recall that $\mathbf{n}_A(a)$ is the unit-normal to A in a pointing to the left. If $c = 0$, C is a single point. Note that a is the point on A closest to x. A special case happens when A is an arc on C, in which case A is a convex arc and C osculates A at every point on A.

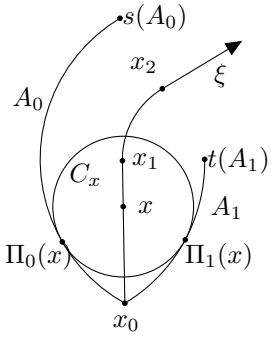

Figure 8

▶ **Observation 14.** *Let A be an arc of ∂P and x an arbitrary point. There exists at most one circle C with center x that osculates A. If such a circle C exists, the radius of C is the distance to the closest point a on A from x and C osculates A in a. If A is a convex arc, then the radius of C is at most the radius of A.*

We denote the bisector curve as $\xi = \xi(A_0, A_1, x_0)$. A_0 and A_1 are arcs of ∂P and x_0 is the point where ξ begins. ξ consists of at most three intervals to be defined below, each of which is part of a conic section. Depending on x_0, ξ might start in the second or third interval. From the context where ξ is used, it will be clear in which interval ξ starts.

Often, but not always, A_0 and A_1 are neighboring arcs and x_0 is their common endpoint, i.e., $x_0 = t(A_0) = s(A_1)$. In this case, x_0 will be a convex vertex of ∂P (as in Figure 8) or A_0 and A_1 are tangential. If A_0 and A_1 are tangential, one of A_0 and A_1 will be a cut arc.

For each point $x \in \xi$, we are going to define a circle C_x and points $\Pi_0(x) \in A_0 \cap C_x$ and $\Pi_1(x) \in A_1 \cap C_x$. The circle C_x has its center at x. The functions Π_0 and Π_1 are the *projections* associated to ξ and C_x is the *clearance circle* associated to ξ at x.

Interval **(1)** consists of points x such that C_x osculates both A_0 and A_1 at points $\Pi_0(x)$ and $\Pi_1(x)$, respectively. Interval **(2)** consists of points x where $s(A_0) = \Pi_0(x)$ and C_x osculates A_1 at a point $\Pi_1(x)$ or $t(A_1) = \Pi_1(x)$ and C_x osculates A_0 at a point $\Pi_0(x)$. Interval **(3)** consists of points x such that $s(A_0) = \Pi_0(x)$ and $t(A_1) = \Pi_1(x)$.

Assume that x_0 is contained in interval **(1)**. In the following, we think of x as a point that traverses ξ from x_0.

Interval (1): ξ consists of the points x such that there is a circle C_x with center x that osculates A_0 and A_1 at points $\Pi_0(x)$ and $\Pi_1(x)$, respectively. Hence, this interval consists of points x on ξ such that the distance to the closest points on A_0 and A_1 is the same. In the general case, as x traverses ξ, $\Pi_0(x)$ moves continuously backward and $\Pi_1(x)$ moves continuously forward. We eventually reach a point x_1 where C_{x_1} contains $s(A_0)$ or $t(A_1)$, and ξ continues with the interval defined in **(2)**. A special case occurs if $t(A_0) = s(A_1)$ and A_0 and A_1 are tangential. Then x moves in the direction $\mathbf{n}_{A_0}(t(A_0))$ until x is the center x_1 of A_0 or A_1. This happens since one of the arcs is a cut arc. Until x reaches x_1, the projections are constant, $\Pi_0(x) = \Pi_1(x) = t(A_0)$. If x_1 is the center of A_0, we define $\Pi_0(x_1) = s(A_0)$. Similarly, if x_1 is the center of A_1, we define $\Pi_1(x_1) = t(A_1)$.

Interval (2): ξ has reached a point x_1 such that C_{x_1} contains $s(A_0)$ or $t(A_1)$. If $s(A_0) = t(A_1)$, ξ stops here. Otherwise, if C_{x_1} contains both $s(A_0)$ and $t(A_1)$, interval **(2)** is degenerate, $x_2 = x_1$, and ξ continues with the interval described in **(3)**. Otherwise, assume without loss of generality that C_x contains $s(A_0)$ and osculates A_1. From here, ξ consists of points x such that there is a circle C_x with center x that contains $s(A_0) = \Pi_0(x)$ and osculates A_1 at a point $\Pi_1(x)$. When x moves along ξ, $\Pi_1(x)$ moves continuously forward. At some point, we reach a point x_2 where C_{x_2} contains $s(A_0)$ and $t(A_1)$, and ξ continues with the interval defined in **(3)**.

Interval (3): At the last point x_2 of interval **(2)**, C_{x_2} contains both of the points $s(A_0)$ and $t(A_1)$. Now, ξ consists of the points x such that there is a circle C_x with center x containing both $\Pi_0(x) = s(A_0)$ and $\Pi_1(x) = t(A_1)$, i.e., ξ follows a half-line. We define the direction of the half-line to be the counterclockwise rotation by an angle of $\pi/2$ of the vector $t(A_1) - s(A_0)$.

When defining the arc $\mathcal{C}$ that specifies the cut we want to perform, we are often searching for the first point x_3 on ξ where C_{x_3} has radius 1. We shall then define $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)]$. Yap [16] showed that each of the intervals **(1)**–**(3)** is a part of a conic section. Using elementary geometry we can decide for each interval in $O(1)$ time if it contains a point x_3 such that C_{x_3} has radius 1.

Note that when $x_0 = t(A_0) = s(A_1)$ and x_0 is a convex vertex of ∂P, then a portion of ξ beginning at x_0 is a subset of the medial axis of ∂P. However, x_3 (as defined above) need not be on the medial axis and can even be outside P, also when $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)]$ becomes a perfect cut arc. See Figure 9 for an example.

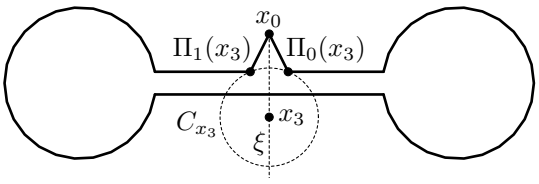

Figure 9

▶ **Observation 15.** *If A_k is a convex arc, $k = 0, 1$, and $x_0 = t(A_0) = s(A_1)$, then for each $a \in A_k$, there exists $x \in \xi$ such that C_x osculates A_k at a, and by Observation 14, the radius of C_x is at most the radius of A_k.*

3.5 Defining the arc $\mathcal{C}$ specifying a cut

The objects in Σ are either **(a)** convex arcs of ∂P_1 with a radius less than 1, **(b)** convex vertices of ∂P_1, or **(c)** cut arcs. Below, we describe how to define $\mathcal{C}$ when the object σ picked by Algorithm 1 is of each of these types.

Case (a): σ is a convex arc with radius less than 1. Consider the curve $\xi = \xi(\sigma, \sigma, x_0)$, where x_0 is the center of σ, which begins in interval **(3)**. Choose the first point $x_3 \in \xi$ such that the associated clearance circle C_{x_3} has radius 1 and define $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)] = C_{x_3}[s(\sigma), t(\sigma)]$. Said in another way, we define $\mathcal{C}$ to be the unit-radius arc counterclockwise from $s(\sigma)$ to $t(\sigma)$ which spans an angle of less than π.

Case (b): σ is a convex vertex of ∂P. Let A_0 be the arc of ∂P before σ and A_1 the arc after σ. Consider the curve $\xi = \xi(A_0, A_1, x_0)$, where $x_0 = \sigma$. Find the first point $x_3 \in \xi$ where C_{x_3} has a radius of 1. If such a point does not exist, we are in the case where $s(A_0) = t(A_1)$, and we let $\mathcal{V}$ be the complete component K of P where σ appears on the boundary (in this case, the boundary of K only has two arcs, A_0 and A_1). Otherwise, we define $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)]$.

Case (c): σ is a cut arc B. We assume that one of the endpoints of B is a convex vertex, since otherwise B is perfect and hence ignored by the algorithm. If one of the endpoints, say $s(B)$, is a concave vertex, then $t(B)$ is a convex vertex, and we do as if $\sigma = t(B)$ as described in case **(b)**. We now assume that none of the endpoints of B are concave vertices. See Figure 10. Let A_0 and A_1 be the arcs preceding and succeeding B on ∂P, respectively. Let $\xi_0 = \xi(A_0, B, s(B))$ and $\xi_1 = \xi(B, A_1, t(B))$. For each $b \in B$, follow the ray starting at b with direction $\mathbf{n}_B(b)$ and define $d_k(b)$ to be the distance from b to the first intersection point with ξ_k, $k = 0, 1$. Since $d_0(s(B)) = 0 < d_1(s(B))$ and $d_1(t(B)) = 0 < d_0(t(B))$, the continuity of ξ_k implies there must be an intersection point x_0 between ξ_0 and ξ_1. By Observation 15, we know that the distance from x_0 to B is at most 1. We now consider the curve $\xi = \xi(A_0, A_1, x_0)$ with associated projections Π_0, Π_1, and clearance circle C_x. The clearance circle C_{x_0} of ξ at x_0 is the same as for ξ_0 and ξ_1. If C_{x_0} osculates A_0 and A_1, ξ begins in interval **(1)**. If C_{x_0} osculates one of A_0, A_1 and contains an endpoint of the other, ξ begins in interval **(2)**. Otherwise, C_{x_0} contains $s(A_0)$ and $t(A_1)$, and ξ begins in interval **(3)**. We find the first point $x_3 \in \xi$ where C_{x_3} has radius 1. If such a point does not exist, we are in the case where $s(A_0) = t(A_1)$, and we let $\mathcal{V}$ be the complete component K of P where σ appears on the boundary. Otherwise, we define $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)]$.

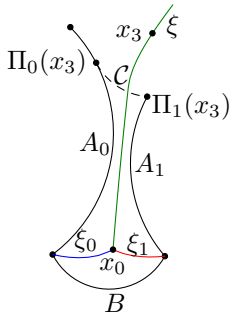

Figure 10

One may object to the algorithm that since a connected component of P can split into more components by a cut, the arcs on Σ are not necessarily well-defined, since subsets of the arc originally added to Σ can be on the boundary of many different connected components of P. However, since the arcs are convex arcs of radius at most 1, the following lemma shows that it is not a problem. The proof will be given in Section 3.7.

▶ **Lemma 16.** *Let A be an arc of ∂P_1 which is a convex arc of any radius or a concave arc with radius at least 1. For any $i \geq 1$, there exists at most one connected component K of P_i such that (a subset of) A is an arc of ∂K.*

3.6 Proof of correctness of Algorithm 1

In this section, we prove that the set $\mathcal{V}$ to be cut away from P is disjoint from every subset of $\overline{P_1}$ of bounded convex curvature. The proof is by contradiction. The idea is that if there is a subset of $\overline{P_1}$ of bounded convex curvature that overlaps $\mathcal{V}$, then Theorem 1 gives the

existence of a unit disk in $\mathcal{V}$ which cannot be there. In general, $\mathcal{V}$ can have a complicated shape and is not easy directly to reason about. Therefore, we describe a superset of $\mathcal{V}$ consisting of *simplices*, which are sufficiently simple that the argument goes through.

A *simplex* S is a closed region and has a boundary ∂S which is a closed, simple curve consisting of two, three, or four circular arcs. One of the arcs D on the boundary is the *door*. The door D has a radius of at most 1 and is clockwise when ∂S is traversed in counterclockwise order. We denote the following as the *simplex condition*, which every simplex is required to satisfy: Let U be the closed disk such that ∂U contains the door D. Then $U \cup S$ does not contain any unit-disk unless U has radius 1, in which case U is the only unit-disk contained in $U \cup S$.

Consider a curve $\xi = (A_0, A_1, x_0)$ with associated projections Π_0, Π_1 and associated clearance circles C_x. For $x_1, x_2 \in \xi$, where x_1 appear before x_2 on ξ, the *simplex candidate* associated to the pair (x_1, x_2) is the curve consisting of $A_0[\Pi_0(x_2), \Pi_0(x_1)]$, $C_{x_1}[\Pi_0(x_1), \Pi_1(x_1)]$, $A_1[\Pi_1(x_1), \Pi_1(x_2)]$, and $C_{x_2}[\Pi_0(x_2), \Pi_1(x_2)]$. The following lemma explains the connection between simplices and simplex candidates.

▶ **Lemma 17.** *Let x_1 and x_2 belong to the same interval of ξ such that x_1 appear before x_2. If there is no $x \in \xi[x_1, x_2)$ for which the clearance circle C_x has a radius of 1 or more, then the simplex candidate associated to (x_1, x_2) is the boundary of a simplex S with door $C_{x_2}[\Pi_0(x_2), \Pi_1(x_2)]$.*

Proof. We prove the lemma when both x_1 and x_2 are in interval **(1)**. The other cases follow from similar arguments. Since x_1 and x_2 are in interval **(1)**, both of the circles C_{x_1} and C_{x_2} osculate A_0 and A_1. Assume that we are not in the case where A_0 and A_1 are tangential, since in that case, the simplex candidate is the boundary of a degenerate simplex and the statement trivial. In the general case, C_{x_1} and C_{x_2} have exactly one point in common with both A_0 and A_1. Therefore, to check that the simplex candidate is simple, we only need to see that $C_{x_1}[\Pi_0(x_1), \Pi_1(x_1)]$ and $C_{x_2}[\Pi_0(x_2), \Pi_1(x_2)]$ do not intersect. If they did, $C_{x_2}[\Pi_0(x_2), \Pi_1(x_2)]$ would cross $C_{x_1}[\Pi_0(x_1), \Pi_1(x_1)$ at some point $a \in C_{x_1}[\Pi_0(x_1), \Pi_1(x_1)$. Since $\Pi_1(x_2)$ is after $\Pi_1(x_1)$ on ∂P, $C_{x_2}[\Pi_0(x_2), \Pi_1(x_2)]$ would have to cross $C_{x_1}[\Pi_0(x_1), \Pi_1(x_1)$ again at some point b. But then $C_{x_2}[b, a]$ is contained in the interior of C_{x_1} and can hence not reach a point $\Pi_1(x_2) \in A_1$.

Let S be the closed region bounded by the simplex candidate. We need to check that S satisfies the simplex condition. Let U be the open disk inside C_{x_2}. Assume for contradiction that there is another open unit-disk D in $S \cup U$. We now move D such that the distance from the center of D to x_2 increases until D reaches the boundary of $S \cup U$. That must happen at an inner point of $A_0[\Pi_0(x_2), \Pi_0(x_1)]$ (resp. $A_1[\Pi_1(x_1), \Pi_1(x_2)]$) since we move D away from U and C_{x_1} has radius less than 1. We then "roll" D forward on A_0 (resp. backward on A_1) until it would not be contained in $S \cup U$ if rolled any further. At this time, it must osculate both A_0 and A_1 at points in $A_0(\Pi_0(x_2), \Pi_0(x_1))$ and $A_1(\Pi_1(x_1), \Pi_1(x_2))$, respectively. But then the center of D is in $\xi(x_1, x_2)$ and D has radius 1, contradicting the assumption. ◀

We say that the simplex S from Lemma 17 is the simplex *associated* to the pair (x_1, x_2). In the following we specify a set $\mathfrak{S}$ of simplices using simplex candidates. The set $\mathfrak{S}$ will be used to prove the optimality of the cut performed by our algorithm when some object σ is handled.

If σ is a convex arc of radius less than 1, $\mathfrak{S}$ consists of one simplex S, which is the simplex associated to the pair (x_0, x_3) as defined in point **(a)** in Section 3.5. Put in another way, S is bounded by the two arcs σ and $\mathcal{C}$, the latter being the door.

If σ is a vertex, consider the curve ξ defined in point **(b)** in Section 3.5. We make one simplex for each of the intervals **(1)**–**(3)** of ξ. We only describe the situation where the center x_3 of $\mathcal{C}$ is in interval **(3)**, since this is the most general case. Let x_1 and x_2 be the points on ξ at the end of interval **(1)** and **(2)**, respectively. We now make one simplex associated to each of the pairs (x_0, x_1), (x_1, x_2), and (x_2, x_3).

If σ is a cut arc B, let A_0, A_1, ξ_0, ξ_1, ξ, x_0, x_3, C_x, Π_0, and Π_1 be as defined in point **(c)** in Section 3.5. Let $\Pi_0^{(k)}$, $\Pi_1^{(k)}$, and $C^{(k)}$ be the projections and clearance circle associated to ξ_k, $k = 0, 1$. The circle C_{x_0} is osculating B at a point q. Note that $C_{x_0}^{(0)} = C_{x_0}^{(1)} = C_{x_0}$, $\Pi_0^{(0)}(x_0) = \Pi_0(x_0)$, $\Pi_1^{(0)}(x_0) = \Pi_0^{(1)}(x_0) = q$, and $\Pi_1^{(1)}(x_0) = \Pi_1(x_0)$.

x_0 must be in interval **(1)** or **(2)** of both ξ_0 and ξ_1, since interval **(3)** consists of points of points y where the circle $C_y^{(k)}$ does not osculate B. If x_0 is in interval **(1)** of ξ_0, we use one simplex associated to the pair $(s(B), x_0)$. If x_0 is in interval **(2)**, let $x_1^{(0)}$ be the end of interval **(1)**. We use simplices associated to the pairs $(s(B), x_1^{(0)})$ and $(x_1^{(0)}, x_0)$. Note that in any case, we make a simplex with the door $C_{x_0}[\Pi_0(x_0), q]$. In a similar way, we make one or two simplices for pairs of points on ξ_1, where one of them has the door $C_{x_0}[q, \Pi_1(x_0)]$.

In the most general situation where x_0 is in interval **(1)**, let x_1 and x_2 be the points on ξ where interval **(1)** and **(2)** ends, respectively. We now make one simplex associated to each of the pairs (x_0, x_1), (x_1, x_2), and (x_2, x_3).

The following is clear from the construction.

▶ **Observation 18.** *For every point $p \in \partial P[s(\mathcal{C}), t(\mathcal{C})]$, there exists a simplex $S \in \mathfrak{S}$ such that $p \in \partial S$. Furthermore, every point of $\partial\mathcal{S}$ is either on $\partial P[s(\mathcal{C}), t(\mathcal{C})]$ or on $\mathcal{C}$, where $\mathcal{S} = \bigcup_{S\in\mathfrak{S}} S$.*

▶ **Lemma 19.** *There exists a set $\mathfrak{S}$ of at most five simplices such that if $\mathcal{S} = \bigcup_{S\in\mathfrak{S}} S$, then $\mathcal{V} \subset \mathcal{S}$ and $\partial\mathcal{S} \subseteq \partial P[s(\mathcal{C}), t(\mathcal{C})] \cup \mathcal{C}$. There is a tree $\mathfrak{T}$ with the nodes being $\mathfrak{S}$ with the following properties. The root of $\mathfrak{T}$ is a simplex with the door $\mathcal{C}$. If a curve $\phi \subset \overline{P}$ exits a simplex $S_0 \in \mathfrak{S}$, it is either through the door of S_0 or through an arc which is the door of a child of S_0 in $\mathfrak{T}$, and in the latter case, ϕ enters the child. If S_0 has a parent in $\mathfrak{T}$ and ϕ exits S_0 through the door of S_0, then ϕ enters the parent.*

Proof. We first see that $\mathcal{V} \subseteq \mathcal{S}$. Let a point $p \in \mathcal{V}$ be given. There exists a curve $\phi \subset P$ from p to a point q on $\partial P[s_\mathcal{C}, t_\mathcal{C}]$ which does not cross $\mathcal{C}$. By Observation 18, q is contained in a simplex of $\mathfrak{S}$ and thus $q \in \mathcal{S}$. We follow ϕ from q to p. When ϕ exits a simplex, it must enter another one by Observation 18 since ϕ does not cross $\partial P[s(\mathcal{C}), t(\mathcal{C})]$ or $\mathcal{C}$. Therefore, $p \in \mathcal{S}$ as well.

It is straightforward to check that the constructed simplices form a tree $\mathfrak{T}$ as described in the lemma. ◀

An example of the tree $\mathfrak{T}$ is given in Figure 11 for case **(c)** from Section 3.5. The following lemma is the heart of our proof of correctness. Recall that the curves in $\mathcal{L}^+$ are open, oriented curves that turn to the left with curvature at most 1.

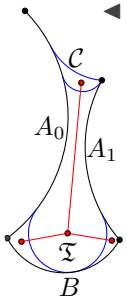

Figure 11

▶ **Lemma 20.** *Let γ be a curve in $\mathcal{L}^+$ contained in $\overline{P}$. Let $\mathcal{C}$ be an arc specifying an area $\mathcal{V}$ to be removed from P and suppose that γ starts at a point $a \in \mathcal{C}$. If γ leaves $\overline{\mathcal{V}}$ again, it is through a point on $\mathcal{C}[s(\mathcal{C}), a]$.*

Proof. Assume for contradiction that γ leaves $\overline{\mathcal{V}}$ through a point not as stated in the lemma. Since the boundary of $\mathcal{V}$ is contained in $\partial P \cup \mathcal{C}$, γ leaves $\overline{\mathcal{V}}$ from a point on $\mathcal{C}(a, t(\mathcal{C})]$. Consider the tree $\mathfrak{T}$ as described in Lemma 19. Let

S be a simplex with maximum distance in $\mathfrak{T}$ to the root such that γ enters S. Let D be the door of S and $s(D)$ and $t(D)$ be the beginning and end of D in counterclockwise order, respectively. Let $\mathcal{D}$ be the open disk such that D is an arc on the circle $\partial\mathcal{D}$. It follows from Lemma 19 that the first time γ enters S, it is through a point on D. Let $e \in D$ be the point closest to $s(D)$ such that γ eventually enters S through e. After e, γ must leave S through a point f on $D(e, t(D))$, since otherwise γ enters the interior of the closed, simple curve defined by $\gamma[a, e]$, $D[s(D), e]$, $\partial P[s(\mathcal{C}), s(D)]$, and $\mathcal{C}[s(\mathcal{C}), a]$, which it cannot exit again by assumption. Hence, Theorem 1 implies that the region bounded by $\partial\mathcal{D}[f, e]$ and $\gamma[e, f]$ contain a unit disk different from $\mathcal{D}$. That contradicts the simplex condition of S. ◀

By the following lemma, we conclude that when the algorithm terminates, $\overline{P}$ is the maximum subset of $\overline{P_1}$ of bounded conve curvature. The lemma follows easily from Lemma 20.

▶ **Lemma 21.** *For every $i = 1, 2, \ldots$, $\overline{P_i}$ contains Q for every set $Q \subset \overline{P_1}$ of bounded convex curvature.*

Proof. Let $Q \subset \overline{P_1}$ be a set of bounded convex curvature. We prove the lemma by showing that each time we perform a cut by removing some set $\mathcal{V}$ from P, $\mathcal{V}$ is disjoint from Q.

Consider for contradiction the first iteration in which Q is not disjoint from $\mathcal{V}$. If $Q \subseteq \mathcal{V}$, then, by Corollary 2, Q contains an open unit-disk D. From the simplex condition, we know that no simplex in $\mathfrak{S}$ contains D. Therefore, we may consider the simplex S with maximum distance to the root that contains a part of ∂D and proceed as in the proof of Lemma 20 to get a contradiction.

If not $Q \subseteq \mathcal{V}$, consider the first point $a \in \mathcal{C}$ where ∂Q enters $\overline{\mathcal{V}}$ when following $\mathcal{C}$ from $s(\mathcal{C})$. ∂Q must exit $\overline{\mathcal{V}}$ through a point on $\mathcal{C}(a, t(\mathcal{C}))$, which contradicts Lemma 20. ◀

3.7 Proof of Lemma 16

The following lemma is necessary to prove Lemma 16.

▶ **Lemma 22.** *The arc $\mathcal{C}$ specifying a cut spans an angle of at most π. Therefore, the same holds for every cut arc of ∂P.*

Proof. Assume that $\mathcal{C}$ spans an angle of more than π. Consider the curve $\xi = \xi(A_0, A_1, x_0)$ and the associated clearance circles C_x and projections Π_0 and Π_1 used to define $\mathcal{C}$. Let x_3 be the center of $\mathcal{C}$. Then $\mathcal{C} = C_{x_3}[\Pi_0(x_3), \Pi_1(x_3)]$. Let x_1 be the first point in the first interval of ξ that contains x_3 (since the intervals overlap at the endpoints, there might be more intervals containing x_3). Note that C_{x_1} has radius less than 1 since otherwise, we would have chosen $x_3 = x_1$. Let S be the simplex associated to the pair (x_1, x_3) (using Lemma 17) and let U be the open unit-disk bounded by C_{x_3}. Since $\mathcal{C}$ spans an angle of more than π, there exists unit-disks in $S \cup U$ different than U, contradicting the simplex condition. ◀

Proof of Lemma 16. Consider the smallest i for which an arc A of ∂P_1 exists such that subsets of A appear as arcs on the boundary of more than one connected component of P_{i+1}. Let $\mathcal{C}$ be the arc specifying the set $\mathcal{V}$ such that $P_{i+1} = P_i \setminus \mathcal{V}$. Since subsets of A appear on more than one component of P_{i+1}, a part of A is removed from the boundary of P_i by the cut. This means that $\mathcal{C}$ intersects A in points a and b when following $\mathcal{C}$ in counterclockwise order, and the part $A(a, b)$ is not on ∂P_{i+1}. A appears on the boundary of three regions R_1, R_2, R_3 of $P_i \setminus \mathcal{C}$, where $A[a, b]$ is an arc on ∂R_2, ∂R_1 contains $s(A)$, and ∂R_3 contains $t(A)$. It could be possible that $R_1 = R_3$, but it does not make a difference for this proof.

Let C be the circle on which $\mathcal{C}$ is an arc. The arc $A(a,b)$ is either in the interior or the exterior of C. Assume first that $A(a,b)$ is in the interior. Note that parts of $\mathcal{C}$ are concave arcs on the boundary of R_1 and R_3. Therefore, if R_2 is the only region among R_1, R_2, R_3 which is removed from P_i, there would be concave cut arcs on ∂P_{i+1}, contradicting Lemma 13.

Consider now the case where $A(a,b)$ is in the exterior of C. We may assume that all arcs of ∂P_1 span an angle of at most π. If an arc spans a larger angle, we can divide it into two arcs each spanning at most π by adding a vertex in the middle of the arc and thus at most doubling the size of the input. Since $\mathcal{C}$ spans an angle of at most π by Lemma 22 and the same holds for A, the radius of A must be less than 1. Assume for contradiction that A is convex. Since R_2 is removed from P_i, a part of $P_i(s(\mathcal{C}), t(\mathcal{C}))$ appear on ∂R_2. $\mathcal{C}[a,b]$ and $A[a,b]$ form a closed, simple curve J. If $\partial P_i(s(\mathcal{C}), t(\mathcal{C}))$ did not cross $\mathcal{C}[a,b]$, R_2 would be the interior of J and R_2 would not be removed from P_1. We can therefore consider the point x where $\partial P_i(s(\mathcal{C}), t(\mathcal{C}))$ enters J for the first time and the point y where it leaves J for the last time. y must be on $A(a,x)$, since P_i is to the left of ∂P_i. If $\partial P_i(s(\mathcal{C}), x)$ is disjoint from $\mathcal{C}$, the curves $\partial P_i[s(\mathcal{C}), x]$ and $\mathcal{C}[s(\mathcal{C}), x]$ form a simple, closed curve that separates $t(\mathcal{C})$ from the points on $\partial P_i[s(\mathcal{C}), x]$ in a neighborhood after y. Therefore, $\partial P_i(y, t(\mathcal{C}))$ must cross $\mathcal{C}(s(\mathcal{C}), a)$ in order to get to $t(\mathcal{C})$. Hence, there are at least five points in $\partial P_1[s(\mathcal{C}), t(\mathcal{C})] \cap \mathcal{C}$. If $\partial P_1(s(\mathcal{C}), x)$ crosses $\mathcal{C}$ somewhere, there are also at least five intersection points. Therefore, $\partial P_1[s(\mathcal{C}), t(\mathcal{C})]$ must consist of at least three arcs, since there are at most two intersection points between two arcs which are not one the same circle. On the other hand, we know from the definition of $\mathcal{C}$ that there can be at most three such arcs. Denote the arcs B_0, B_1, B_2 in the order in which they appear on $\partial P_1[s(\mathcal{C}), t(\mathcal{C})]$, where $B_1 = \sigma$ is the cut arc that was picked from Σ. Assume first that $x \in B_0$. Since either B_0 and B_1 are tangential, or $s(B_1)$ is a convex vertex, it is clear that y has to be on $\mathcal{C}(x,b)$, which is a contradiction. It is also impossible that $x \in B_2$, since in that case, B_2 should have four intersection points with $\mathcal{C}$.

Assume then that $x \in B_1$. Since B_1 and $\mathcal{C}$ both have radius 1, the angle of the arc on B_1 from a point where it exits C to the point where it enters C would be more than π. Hence, B_1 does not leave the interior of J again. Then B_2 must have three intersection points with $\mathcal{C}$, which is a contradiction.

Hence, if parts of A appear on the boundary of more than one component of P, A must be a concave arc with radius lass than 1. ◀

3.8 Proof that Algorithm 1 performs a linear number of iterations

Since the set $\mathcal{V}$ that we remove from P is contained in one connected component of P, the algorithm only makes changes to one connected component of P at a time. A vertex v of ∂P is *improper* if v is a convex vertex and v is not a vertex of ∂P_1. An improper vertex can be introduced in two different ways: 1) If the arc $\mathcal{C}$ specifying a cut intersects a cut arc and the intersection point v is on $\partial\mathcal{V}$, then v becomes a convex and hence improper vertex. 2) If the arc $\mathcal{C}$ begins or ends at the endpoint $v \notin \partial P_1$ of a cut arc which is not a convex vertex, v might be a convex and hence improper vertex after the cut has been performed.

Consider the first time an improper vertex appears. Let v be the vertex and let K be the connected component such that $v \in \partial K$. Then v will be removed by the subsequent cut performed in K since a cut arc with an endpoint at v is closest to the top of Σ among all objects in $\Sigma \cap \partial K$. It follows inductively that there can be one or two improper vertices in each component, and after the next iteration performed in the component, none of them will be left.

A vertex v of ∂P which is not a point on ∂P_1 (note the difference between being a *vertex*

of ∂P_1 and being a *point* on ∂P_1) can only be created if the arc $\mathcal{C}$ specifying a cut to be performed in P intersects a cut arc at v. It follows that v is a convex and hence improper vertex. Equivalently, every proper vertex is on ∂P_1.

Consider an endpoint c of the arc $\mathcal{C}$ specifying a cut to be performed in P_i, where $P_{i+1} = P_i \setminus \mathcal{V}$. If c is an inner point of an arc A of ∂P_i, Lemma 13 together with Observation 14 gives that A is not a cut arc since the circle containing $\mathcal{C}$ osculates A at c. Hence, A is (a subset of) an arc of ∂P_1 and $c \in \partial P_1$. Furthermore, due to the way $\mathcal{C}$ is defined, an improper vertex of ∂P_i will never be chosen as c. Also in this case, we get that c is on ∂P_1. We summerize these observations in the following lemma.

▶ **Lemma 23.** ***1.*** *There are at most two improper vertices of ∂K for each connected component K of P. If there are two, they are endpoints of the same cut arc.*
2. *All improper vertices of ∂K, where K is a connected component of P, are removed by the following cut performed in K.*
3. *The endpoints $s(\mathcal{C})$ and $t(\mathcal{C})$ of an arc $\mathcal{C}$ specifying a cut are both points on ∂P_1.*

▶ **Lemma 24.** *The algorithm performs $O(n)$ iterations.*

Proof. We divide the iterations into the types 1–5 defined below and prove that there is a linear number of each type. We bound the iterations of type 1 by observing that each such iteration corresponds to introducing an edge in a planar multigraph with $2n$ vertices, where each pair of vertices can be connected by at most $O(1)$ edges. Hence, Euler's formula implies that $O(n)$ edges are created in total. The iterations of types 2–5 can be bounded by the number of iterations of type 1 or by observing that each iteration completely removes one of the objects of ∂P_1 from ∂P. Hence there are $O(n)$ of each type.

Consider what happens during iterations i when the algorithm picks an object σ from Σ. Let $\mathcal{V}$ be the set removed from P_i during the iteration, i.e., $P_{i+1} = P_i \setminus \mathcal{V}$, where possibly $\mathcal{V} = \emptyset$. Lemma 23 implies that the iteration is of one of the following types.

1. $\mathcal{V}$ is specified by an arc $\mathcal{C}$ and each endpoint of $\mathcal{C}$ is either a vertex of ∂P_1 or an inner point of an arc of ∂P_1 which is not a vertex of ∂P_i.
2. $\mathcal{V}$ is specified by an arc $\mathcal{C}$ and one or both endpoints of $\mathcal{C}$ is a vertex v of ∂P_i which is not a vertex of ∂P_1. Hence, such a vertex v is the endpoint of a cut arc of ∂P_i.
3. A complete connected component K is removed from P, i.e., $\mathcal{V} = K$.
4. $\mathcal{V} = \emptyset$ and nothing happens since σ is not anymore on ∂P.
5. $\mathcal{V} = \emptyset$ and nothing happens since σ is a perfect cut arc.

Consider cuts of type 1. We see that each such cut corresponds to adding an edge to a planar graph G with $2n$ nodes. There is one node in G for each arc and vertex of ∂P_1. We choose a curve ϕ from $s(\mathcal{C})$ to $t(\mathcal{C})$ contained in $\mathcal{V}$ as the edge corresponding to the cut. This is possible since $\mathcal{V}$ is connected by Lemma 12. A vertex v of ∂P_1 is incident to ϕ if v is an endpoint of ϕ. An arc A of ∂P_1 is incident to ϕ if an endpoint of ϕ is an inner point of A. Let A_0 and A_1 be the objects of ∂P_1 incident to ϕ such that $s(\mathcal{C}) \in A_0$ and $t(\mathcal{C}) \in A_1$. Let C be the circle containing $\mathcal{C}$. Observe that if A_0 (resp. A_1) is an arc, then C osculates A_0 at $s(\mathcal{C})$ (resp. A_1 at $t(\mathcal{C})$).

We say that A_0 and A_1 are *parallel* if:

- A_0 and A_1 are concentric arcs, one is concave, the other is convex, and the radius of the convex is 2 larger than the radius of the concave.
- A_0 and A_1 are two parallel line segments (in the ordinary sense) with opposite directions and distance 2.
- A_1 (resp. A_0) is a vertex while A_0 (resp. A_1) is convex arc with radius 2 and center A_1 (resp. A_0).

If A_0 and A_1 are parallel, there are infinitely many unit-circles that osculates A_0 or A_1 like C does. Otherwise, there are only $O(1)$. Hence, by Euler's formula, we are adding $O(n)$ edges to G that connect non-parallel objects.

We now verify that even if A_0 and A_1 are parallel, we make $O(1)$ edges between them in G. Consider the first case where A_0 and A_1 are concentric, A_0 is convex, A_1 is concave, and the radius of A_0 is 2 larger than the radius of A_1. The other types of parallel arcs are handled in a similar way. We consider the case where more than one cut arc is made during the execution of the algorithm which starts on A_0, ends on A_1, and is tangential to both arcs. Assume that the first such arc is $\mathcal{C}_1$ from $s(\mathcal{C}_1) \in A_0$ to $t(\mathcal{C}_1) \in A_1$ created during iteration i. Consider the next time where such a cut arc $\mathcal{C}_2$ is made from $s(\mathcal{C}_2) \in A_0$ to $t(\mathcal{C}_2) \in A_1$ during iteration k, $k > i$. The boundary interval $\partial P_{i+1}[s(\mathcal{C}_2), t(\mathcal{C}_2)]$ contains no convex vertices or convex arcs of radius less than 1. Therefore, some cut arc must have introduced a convex vertex on $\partial P(s(\mathcal{C}_2), t(\mathcal{C}_2))$. Hence, during some iteration l, where $i < l < k$, a cut arc B is introduced from $x \in \partial P_{i+1}(s(\mathcal{C}_2), t(\mathcal{C}_2))$ to $y \in \partial P_{i+1}(x, t(\mathcal{C}_2))$, where $\mathcal{B}(x, y) \subset P_{i+1}$. From Lemma 13 we know that B is convex when traversed from x to y. But such an arc B does not exists for any $x \in P_{i+1}(s(\mathcal{C}_2), t(\mathcal{C}_2))$ due to the assumptions on A_0 and A_1. Therefore, neither B nor $\mathcal{C}_2$ can appear as cut arcs and the number of edges between A_0 and A_1 in G is $O(1)$.

In a similar way, we see that the total number of perfect cut arcs created while executing the algorithm is $O(n)$ by describing a planar graph H. The vertex set of H is also the vertices and arcs of ∂P_1. If a cut creates a new cut arc $\mathcal{C}' \subseteq \mathcal{C}$ of ∂P_{i+1} which is perfect, we add the edge $\mathcal{C}'$ to H. An endpoint of $\mathcal{C}'$ cannot be an interior point of a cut arc since $\mathcal{C}'$ is perfect. Therefore, $s(\mathcal{C}'), t(\mathcal{C}') \in P_1$. Hence, we can bound the number of edges in H in a similar way as the number of edges in G.

We say that a cut arc is *active* if it is in Σ. We see inductively that there can be at most one active cut arc in each connected component of P. Namely, assume that there is one active cut arc $\mathcal{C}$ in the component K of P. Assume that iteration i is the following iteration performed in K. We must have $\sigma = \mathcal{C}$, since $\mathcal{C}$ was added to the top of Σ when it was created. During the iteration, $\mathcal{C}$ either disappears from ∂P when removing $\mathcal{V}$ or no change to P is made since $\mathcal{C}$ is perfect. This means that all inactive cut arcs are perfect.

Consider now a cut of type 2. Assume that an endpoint of $\mathcal{C}$ is an endpoint of a cut arc A of ∂P_i. The cut arc A is inactive and hence perfect. After the cut, v is a convex vertex. After one more iteration in the component containing A, A will not anymore be an arc of ∂P. Hence, there are at most 2 iterations of type 2 for any perfect arc, so there are $O(n)$ iterations of type 2.

Consider a cut of type 3. Let K be the component removed by the cut. Note that there are at most three arcs on ∂K. Assume first that there are at least two cut arcs on ∂K. One of them is not active and is hence perfect. Therefore, there can be removed at most $O(n)$ such components K. Assume now that there is at most one cut arc on ∂K. If there are three arcs on ∂K, there is also a vertex of ∂P_1. It can happen at most n times that such a component is removed. The only remaining case is that ∂K consists of one cut arc $\mathcal{C}$ and an arc A' which is a subset of an arc A of ∂P_1. A' must be a convex arc, a line segment, or a concave arc with radius more than 1. Therefore, by Lemma 16, there is no other connected component of P where another subset of A is an arc. Hence, there can be removed at most n such components K.

Consider an iteration of type 4. If σ is not anymore on ∂P, σ must be an object of ∂P_1, since every active cut arc is the uppermost object in Σ among all objects of the same connected component of P as σ. Therefore, there are at most $2n$ iterations of this type.

Consider the iterations of type 5. As we saw before, there are made $O(n)$ perfect cut arcs during the execution of the algorithm, which bounds the number of this type. ◀

Since every cut arc is eventually picked from Σ by Algorithm 1, the linear bound on the number of iterations leads to the following lemma.

▶ **Lemma 25.** *The total number of cut arcs created while running Algorithm 1 is $O(n)$. Likewise, the total number of different vertices appearing on ∂P while running Algorithm 1 is $O(n)$.*

We have now proved the following theorem.

▶ **Theorem 26.** *Given a simply-connected open region P_1 bounded by a curvilinear polygon consisting of n line segments and circular arcs, there is a unique maximum set $Q \subseteq \overline{P_1}$ of bounded convex curvature which contains every set $Q' \subseteq \overline{P_1}$ of bounded convex curvature. ∂Q consists of $O(n)$ line segments and circular arcs, and Algorithm 1 computes Q in $O(n)$ iterations.*

4 Implementation

In this section, we show how the algorithm can be implemented so that it uses $O(n \log n)$ time in total and $O(n)$ space.

Let $\mathcal{C}$ be the arc specifying a cut to be performed in P. Let $R_1, \ldots, R_r$ be the connected components of $P \setminus \mathcal{C}$, and assume without loss of generality that the sets $R_1, \ldots, R_s$, $s \leq r$, have to be removed since their boundaries contain a part of $\partial P(s(\mathcal{C}), t(\mathcal{C}))$. For ease of presentation we assume that $\mathcal{C}$ contains no vertices of ∂P except for possibly the endpoints $s(\mathcal{C})$ and $t(\mathcal{C})$. Algorithm 2 traverses the boundary of each of the regions $R_1, \ldots, R_s$ and introduces the new cut arcs in order to remove the regions from P.

Algorithm 2: `PerformCut`$(\mathcal{C})$

```
1  U ← {s(C)}.
2  while U ≠ ∅
3      Remove a point e from U. If e is not a vertex of ∂P, create a vertex at e.
4      a ← e.
5      repeat
6          Let b ← TraverseP(C, a, e). If b is not a vertex of ∂P, create a vertex at b.
7          Mark all objects on ∂P[a, b] as removed, except for vertices a and b.
8          If b ∈ ∂P(s(C), t(C)), add b to U.
9          if C enters or leaves P at b
10             a ← TraverseC(C, b). If a is not a vertex of ∂P, create a vertex at a.
11             Change ∂P by setting the arc succeeding a and preceding b to C[a, b].
12         else
13             a ← e.
14     until a = e
```

The algorithm traverses $\partial P(s(\mathcal{C}), t(\mathcal{C}))$ from $s(\mathcal{C})$ and removes the regions $R_1, \ldots, R_s$ one by one. Each point in the set U is on the boundary of one of the regions $R_1, \ldots, R_s$ which has not so far been removed. The loop at line 5 traverses one such region.

The subroutine `TraverseP`$(\mathcal{C}, a, e)$ traverses ∂P from a until we meet e or a point where $\mathcal{C}$ enters or exits P. The point b where we stop traversing is returned.

The subroutine `TraverseC`$(\mathcal{C}, b)$ follows $\mathcal{C}$ from b through P until it exits P at some point a, which is returned. Since we assumed that no vertex of ∂P is an inner point of $\mathcal{C}$, it is uniquely defined which direction of $\mathcal{C}$ to follow.

Since P is simply-connected, the regions $R_1, \ldots, R_r$ induce a tree T such that there is an edge in T between two regions R_j and R_k if there is an arc $\mathcal{C}(x, y) \subset \partial R_j \cap \partial R_k$ for $x \neq y$. By Lemma 12, $\mathcal{V}$ is connected, so the regions $R_1, \ldots, R_s$ induce a subtree in T. Therefore, each of the regions in T is traversed at least once by Algorithm 2. On the other hand, a region is removed from P while its boundary is traversed, so each region is traversed at most once. Hence, we have the following lemma.

▶ **Lemma 27.** *Algorithm 2 traverses the boundary of each region $R_1, \ldots, R_s$ exactly once and thus correctly computes $P \leftarrow P \setminus \mathcal{V}$ as specified by the arc $\mathcal{C}$.*

The point `TraverseC`$(\mathcal{C}, b)$ is the first intersection point between $\mathcal{C}$ and ∂P when following $\mathcal{C}$ from b in the relevant direction. Cheng et al. [5] describes an efficient solution to the following problem. For a simple polygon V, preprocess V such that queries of the following kind can be answered efficiently: Given a circular unit-radius arc beginning at some point in the interior of V, find the first intersection point between the arc and V if it exists. The algorithm requires $O(n)$ space, uses $O(n \log n)$ time on preprocessing, and answers queries in $O(\log n)$ time, where n is the number of vertices of V.

It is straightforward to generalize the method of Cheng et al. to curvilinear polygons. We apply the preprocessing to the original input ∂P_1. Thus, by querying an arc following $\mathcal{C}$ from b in the direction through P, we know the point where $\mathcal{C}$ exits P_1. However, the arc $\mathcal{C}$ can exit P before it exits P_1, namely if and only if it crosses a cut arc of ∂P. In the following, we show how to detect if that is the case or not.

Since $\mathcal{C}$ enters or exits P at b, there exists a point $d \in \mathcal{C} \cap P$ such that $\mathcal{C}(d, b) \subset P$. In the following, d denotes such a point. It is not necessary to compute d, it is merely a tool in our analysis. We know from Lemma 23 that $s(\mathcal{C}), t(\mathcal{C}) \in \partial P_1$. Using the circular ray shooting data structure, we can find a closed subset $\mathcal{C}' \subseteq \mathcal{C}$ such that $s(\mathcal{C}'), t(\mathcal{C}') \in \partial P_1$, $\mathcal{C}'(s(\mathcal{C}'), t(\mathcal{C}')) \subset P_1$, and $b \in \mathcal{C}'$. Hence also $d \in \mathcal{C}'$. The following lemma says that if the arc $\mathcal{C}'$, when traversed in any direction from d, enters a removed area, i.e., a connected component of $P_1 \setminus P$, then it stays in that removed area.

▶ **Lemma 28.** *Let x be one of the endpoints of $\mathcal{C}'$. There is at most one cut arc A of ∂P which intersects $\mathcal{C}'(d, x)$. There is such an arc A if and only if $x \notin \partial P$. If A exists, $\mathcal{C}'(d, x)$ and A intersect at a unique point y, and $\mathcal{C}'[y, x) \subset P_1 \setminus P$.*

Proof. Assume first that there is a cut arc A of ∂P intersecting $\mathcal{C}'(d, x)$ at a point y. Assume for contradiction that $\mathcal{C}'(y, x)$ intersects a cut arc of ∂P at a point z. $\mathcal{C}'$ exits P at y and enters P at z. We may consider the smallest i such that $\mathcal{C}'(d, x)$ contains an interval $\mathcal{C}'[y, z]$ which is outside P_{i+1}. Thus, $\mathcal{C}'[y, z] \subset \mathcal{V}$, where $P_{i+1} = P_i \setminus \mathcal{V}$. Let B be the arc that specifies the set $\mathcal{V}$. Consider the intersection point $w \in B \cap \mathcal{C}'(d, x)$ where $\mathcal{C}'(d, x)$ enters or leaves $\mathcal{V}$ which is closest to $s(B)$. Note that $\mathcal{C}'$ is a curve in $\mathcal{L}^+$ independently of the direction in which we traverse $\mathcal{C}'$. Therefore, Lemma 20 gives that when following $\mathcal{C}'$ into $\mathcal{V}$ from w, the point where it exits $\mathcal{V}$ is on $B[s(B), w)$. That is a contradiction since w was as close to $s(B)$ as possible. Hence, from y an on, $\mathcal{C}'$ is contained in $P_1 \setminus P$ and $x \notin \partial P$.

On the other hand, if $\mathcal{C}'(d, x)$ does not intersect any cut arc of ∂P, then $\mathcal{C}(d, x)$ is contained in P and therefore $x \in \partial P$. ◀

It is possible that $\mathcal{C}'(s(\mathcal{C}'), d)$ and $\mathcal{C}'(d, t(\mathcal{C}'))$ intersect the same cut arc B. That is easy to detect since we are given a point $b \in \partial P \cap \mathcal{C}'$ where $\mathcal{C}'$ enters of leaves P. Thus, we check if b is on a cut arc B and if $\mathcal{C}(b, x)$ intersects the same arc. Otherwise, every cut arc A satisfies that $A \cap \mathcal{C}'(s(\mathcal{C}), d) = \emptyset$ or $A \cap \mathcal{C}'(d, t(\mathcal{C})) = \emptyset$, and we have the following lemma.

▶ **Lemma 29.** *Let A be a cut arc of ∂P such that $s(A), t(A) \in \partial P_1$. Assume that $A \cap \mathcal{C}'(s(\mathcal{C}'), d) = \emptyset$ or $A \cap \mathcal{C}'(d, t(\mathcal{C}')) = \emptyset$. Then, the following two statements hold:*

- *A and $\mathcal{C}'(s(\mathcal{C}'), d)$ intersect if and only if the endpoints of $\mathcal{C}'$ and A appear in the order $s(\mathcal{C}'), t(A), t(\mathcal{C}'), s(A)$ on ∂P_1.*
- *A and $\mathcal{C}'(d, t(\mathcal{C}'))$ intersect if and only if the endpoints of $\mathcal{C}'$ and A appear in the order $s(\mathcal{C}'), s(A), t(\mathcal{C}'), t(A)$ on ∂P_1.*

Proof. We prove the first statement. The second follows from similar arguments. Assume that A and $\mathcal{C}'(s(\mathcal{C}'), d)$ intersect at a point a. By Lemma 28, we know that $\mathcal{C}'(s(\mathcal{C}'), a)$ does not intersect A. Since $\mathcal{C}'(b, t(\mathcal{C}'))$ does not intersect A by assumption, the endpoints of $\mathcal{C}'$ are on different sides of A in P_1. A divides P_1 into two regions, one with a boundary including $\partial P_1(s(A), t(A))$ and the other with a boundary including $\partial P_1(t(A), s(A))$. The portion $\mathcal{C}'(s(\mathcal{C}'), a]$ is outside P. Hence $\mathcal{C}'$ enters P at a, so Lemma 13 gives that $s(\mathcal{C}') \in \partial P_1(s(A), t(A))$ and $t(\mathcal{C}') \in \partial P_1(t(A), s(A))$, or, equivalently, that the order of the points are as stated in the lemma.

If the order of the endpoints is as stated, the endpoints of $\mathcal{C}'$ are on different sides of A, and hence $\mathcal{C}'$ and A intersect at some point a. Lemma 28 gives that $\mathcal{C}'(s(\mathcal{C}'), a] \subset P_1 \setminus P$. Hence, $a \in \mathcal{C}'(s(\mathcal{C}'), d)$. ◀

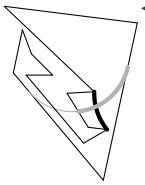

Figure 12

See Figure 12, which illustrates Lemma 29. The black arc A is the only cut arc of ∂P. The thick part of the grey arc $\mathcal{C}$ is $\mathcal{C}'$.

Lemma 29 leads to our method for finding cut arcs intersecting an arc $\mathcal{C}'$ by searching after arcs with endpoints on specific portions of ∂P_1. We associate to each point p on ∂P_1 a unique number $\varphi(p) \in [0, n)$. Let the vertices of ∂P_1 be $v_0, v_1, \ldots, v_{n-1}, v_n$, where $v_0 = v_n$. We set $\varphi(v_i) = i$ for $i < n$. For the points p on the arc between two vertices v_i and v_{i+1}, we interpolate between i and $i+1$ to uniquely define $\varphi(p)$. For an arc A with $s(A), t(A) \in \partial P_1$, we define an associated point $\theta(A) \in [0, n) \times (0, 2n)$ in the following way:

$$\theta(A) = \begin{cases} (\varphi(s(A)), \varphi(t(A))) & \text{if } \varphi(s(A)) < \varphi(t(A)) \\ (\varphi(s(A)), \varphi(t(A)) + n) & \text{otherwise.} \end{cases}$$

Lemma 29 then leads to the following.

▶ **Lemma 30.** *Let A be a cut arc of ∂P such that $s(A), t(A) \in \partial P_1$. Assume that $A \cap \mathcal{C}'(s(\mathcal{C}'), d) = \emptyset$ or $A \cap \mathcal{C}'(d, t(\mathcal{C}')) = \emptyset$. Let $\theta(\mathcal{C}') = (x, y)$.*

- *A and $\mathcal{C}'(s(\mathcal{C}'), d)$ intersect if and only if*
  - *$y < n$ and $\theta(A)$ is in $[0, x) \times (x, y)$ or $(y, n) \times (x + n, y + n)$, or*
  - *$y \geq n$ and $\theta(A)$ is in $(y - n, x) \times (x, y)$.*
- *A and $\mathcal{C}'(d, t(\mathcal{C}'))$ intersect if and only if*
  - *$y < n$ and $\theta(A)$ is in $(x, y) \times (y, x + n)$, or*
  - *$y \geq n$ and $\theta(A)$ is in $[0, y - n) \times (y - n, x)$ or $(x, n) \times (y, x + n)$.*

Proof. Lemma 29 says that A and $\mathcal{C}'(s(\mathcal{C}'), d)$ intersect if and only if the endpoints of $\mathcal{C}'$ and A appear in the order $s(A), s(\mathcal{C}'), t(A), t(\mathcal{C}')$ on ∂P_1. If we assume that $y < n$ and $\varphi(s(A)) < \varphi(t(A))$, this is equivalent to $0 \leq \varphi(s(A)) < \varphi(s(\mathcal{C}')) < \varphi(t(A)) < \varphi(t(\mathcal{C}')) < n$, which means that $\theta(A) \in [0, x) \times (x, y)$. The other cases follows in a similar way. ◀

For each cut arc A of ∂P where $s(A), t(A) \in \partial P_1$, we store the point $\theta(A)$ in a data structure Θ. It is necessary to add new points to and delete points from Θ as the algorithm proceeds, since new cut arcs are created and other no longer appear on ∂P. We need to be able to find a point $\theta(A)$ in a rectangle specified by $\mathcal{C}'$ as stated in Lemma 30. Therefore, we implement Θ as a fully dynamic orthogonal range searching structure as described by Blelloch [2]. Algorithm 3 sketches how to implement `TraverseC`.

Algorithm 3: `TraverseC`$(\mathcal{C}, b)$

1 Use the circular ray shooting data structure to find the arc $\mathcal{C}' \subseteq \mathcal{C}$ such that $\mathcal{C}'(s(\mathcal{C}'), t(\mathcal{C}')) \subset P_1$, $s(\mathcal{C}'), t(\mathcal{C}') \in \partial P_1$, and $b \in \mathcal{C}'$.
2 Let $z \in \{s(\mathcal{C}'), t(\mathcal{C}')\}$ be the endpoint of $\mathcal{C}'$ when following $\mathcal{C}'$ from b through P.
3 If b is on a cut arc B of ∂P and $\mathcal{C}(b, z)$ intersects B at a point a, return a.
4 Ask the data structure Θ if any point $\theta(A)$ is in the rectangle(s) as specified by Lemma 30. If so, return the intersection point between A and $\mathcal{C}'(b, z)$.
5 Return z.

It is now possible to bound the running time and memory requirement of Algorithm 1 when using our suggested implementation.

▶ **Theorem 31.** *Algorithm 1 can be implemented so that it runs in time $O(n \log n)$ and uses $O(n)$ space.*

Proof. We first bound the space and time used on maintaining and querying the data structure Θ. From Lemma 25, we know that there are $O(n)$ points in Θ at any time. Blelloch [2] describes how to implement Θ using $O(n)$ space so that insertions and deletions are performed in $O(\log n)$ amortized time and orthogonal range reporting queries in $O\left(\log n + k\frac{\log n}{\log\log n}\right)$ time, where k is the number of reported points. In our case, due to Lemma 28, there is 0 or 1 point to report in each query. Therefore, we use $O(n \log n)$ time on Θ in total.

For each cut arc of ∂P, we perform 2 circular ray shooting queries. Thus, we use $O(n \log n)$ time on that in total.

The portion of ∂P we traverse when calling `TraverseP`$(\mathcal{C}, a, e)$ is removed from P and hence never traversed again. Therefore, the time used on `TraverseP` is bounded by the number of vertices occuring on ∂P during the execution of the algorithm. By Lemma 25, we use $O(n)$ time on `TraverseP` in total. ◀

Acknowledgements

We want to thank Jacob Holm, Ralph Howard, Mathias Bæk Tejs Knudsen, and Eva Rotenberg for useful suggestions and discussions

References

1 H.-K. Ahn, O. Cheong, J. Matoušek, and A. Vigneron. Reachability by paths of bounded curvature in a convex polygon. *Computational Geometry*, 45(1):21–32, 2012.
2 G.E. Blelloch. Space-efficient dynamic orthogonal point location, segment intersection, and range reporting. In *Proceedings of SODA*, pages 894–903, 2008.
3 F. Bonahon. *Low-dimensional geometry: From Euclidean surfaces to hyperbolic knots*, volume 49 of *Student Mathematical Library*. American Mathematical Society, 2009.

4 A.M. Bruckner. *Differentiation of real functions*, volume 659 of *Lecture Notes in Mathematics*. Springer-Verlag, 1978.

5 S.-W. Cheng, O. Cheong, H. Everett, and R. van Oostrum. Hierarchical decompositions and circular ray shooting in simple polygons. *Discrete and Computational Geometry*, 32:401–415, 2004.

6 H.S. Choy and K.W. Chan. A corner-looping based tool path for pocket milling. *Computer-Aided Design*, 35(2):155–166, 2003.

7 X. Han and L. Tang. Precise prediction of forces in milling circular corners. *International Journal of Machine Tools and Manufacture*, 88:184–193, 2015.

8 M. Held. Voronoi diagrams and offset curves of curvilinear polygons. *Computer-Aided Design*, 30(4):287–300, 1998.

9 R. Howard and A. Treibergs. A reverse isoperimetric inequality, stability and extremal theorems for plane-curves with bounded curvature. *Rocky Mountain Journal of Mathematics*, 25(2):635–684, 1995.

10 H. Iwabe, Y. Fujii, K. Saito, and T. Kishinami. Study on corner cut by end mill. *International Journal of the Japan Society for Precision Engineering*, 28(3):218–223, 1994.

11 J. McCleary. *Geometry from a differentiable viewpoint*. Cambridge University Press, 2012.

12 S.C. Park and Y.C. Chung. Mitered offset for profile machining. *Computer-Aided Design*, 35(5):501–505, 2003.

13 V. Pateloup, E. Duc, and P. Ray. Corner optimization for pocket machining. *International Journal of Machine Tools and Manufacture*, 44(12):1343–1353, 2004.

14 G. Pestov and V. Ionin. The largest possible circle imbedded in a given closed curve. *Doklady Akademii Nauk SSSR*, 127(6):1170–1172, 1959.

15 H. Rohnert. Shortest paths in the plane with convex polygonal obstacles. *Information Processing Letters*, 23(2):71–76, 1986.

16 C.K. Yap. An $O(n \log n)$ algorithm for the voronoi diagram of a set of simple curve segments. *Discrete and Computational Geometry*, 2(1):365–393, 1987.

17 Z.Y. Zhao, C.Y. Wang, H.M. Zhou, and Z. Qin. Pocketing toolpath optimization for sharp corners. *Journal of Materials Processing Technology*, 192:175–180, 2007.